UNIVERSITÀ DEGLI STUDI DI MILANO

FACOLTÀ DI SCIENZE E TECNOLOGIE

Corso di Dottorato in Matematica

Semistable models of hyperelliptic curves in the wild case

&

Differential operators on p-adic modular forms

Supervisor:
Prof. Fabrizio ANDREATTA

PhD thesis by:
Leonardo FIORE

February 4th, 2024

Contents

Preface

This PhD thesis is divided into two parts, collecting two independent pieces of work I completed during my PhD, both in the realm of number theory and arithmetic geometry.

The first part presents a joint work with J. Yelton, regarding semistable models of hyperelliptic curves in the wild case. To this subject, which was already the topic of my MSc. thesis, I devoted part of my research work during my PhD, until completing it during summer 2022.

The second part of this thesis discusses some research findings related to a completely unrelated topic proposed to me by my PhD advisor Fabrizio Andreatta, namely the one of geometric constructions of differential operators on sheaves of p-adic modular forms.

Ringraziamenti

Vorrei ringraziare, anzitutto, Ananyo Kazi, con cui ho condiviso il percorso del dottorato: è stato per me opportunità di confronto e di compagnia, e a lui rivolgo i miei auguri perché possa trovare soddisfazione nel percorso accademico che ha scelto, diversamente da me, di proseguire.

Una parte del lavoro di dottorato è consistita nella continuazione della collaborazione con Jeffrey Yelton. Ricordo con il piacere e il gusto del viaggio il soggiorno ad Atlanta in occasione del quale abbiamo completato un articolo lungo e non poco travagliato.

È tuttavia con il mio relatore Fabrizio Andreatta che ho il debito di gratitudine più grande, sotto ogni rispetto. Ma, sopra tutto il resto, desidero ringraziarlo per la disponibilità, la simpatia, i toni franchi e il garbo cordiale che ha saputo donare alla nostra collaborazione; in una parola, per l'esempio umano che ha per me significato, e in cui riconosco il valore più grande, e più durevole, che l'esperienza accademica appena conclusa abbia per me avuto.

Ringrazio infine i revisori per i loro commenti, e per i rilievi molto pertinenti che mi hanno consentito di migliorare, in più punti, l'esposizione della materia.

Part 1

Arithmetic of hyperelliptic curves

CHAPTER 1

Introduction

The focus of this work is to investigate the reduction types of hyperelliptic curves over discrete valuation fields. Given a complete discrete valuation field K of characteristic different from 2 with algebraically closed residue field, our starting point is to consider a *hyperelliptic curve* Y over K; that is, Y/K is a smooth projective curve of positive genus admitting a degree-2 morphism onto the projective line $\mathbb{P}^1_K$.

This work is concerned with constructing a semistable model of a given hyperelliptic curve Y/K and understanding the structure of the special fiber of a semistable model of Y. As this problem is already entirely understood in the case that the residue characteristic is not 2 and the procedure in that case can be described entirely in terms of the distances between the branch points with respect to the p-adic metric on K, our primary focus will be on the case where the residue characteristic is 2. The increased complexity of the problem for this case arises from the fact that a hyperelliptic curve comes with a degree-2 map to the projective line: the fact that this degree is the same as the residue characteristic implies that we are in a "wild setting". Problems involving reduction of curves in the "wild case", in which one studies semistable models of curves with a degree-p map to the projective line over residue characteristic p, have been investigated in a number of works in recent decades (see §1.3 below), but mainly in the situation where the branch points of the map $Y \to \mathbb{P}^1_K$ are p-adically equidistant. In this work, we will consider general hyperelliptic curves over residue characteristic 2, with a particular focus on the relationship between the combinatorial data of how the branch points are "clustered" and the structure of the special fiber of a semistable model.

1. Our main problem

It is well known that an affine chart for a hyperelliptic curve Y/K of genus $g \geq 1$ is given by an equation of the form

$$y^2 = f(x) = c\prod_{i=1}^{d}(x - a_i), \tag{1}$$

where $f(x) \in K[x]$ is a polynomial of degree $d \in \{2g+1, 2g+2\}$ that does not have multiple roots, $c \in K^\times$ is the leading coefficient of $f(x)$, and the elements $a_i \in \bar{K}$ are the roots of f. We call f the *defining polynomial* of (this chart of) the hyperelliptic curve Y. The degree-2 morphism of Y onto the projective line is given simply by the coordinate function x; this morphism is branched precisely at each of the roots of f as well as, in the case that $d = 2g + 1$ (in other words, when f has odd degree), at the point ∞. After applying an appropriate automorphism of the projective line (i.e. a suitable change of coordinate) which moves one of the branch points to ∞, we obtain an equation of the form in (1) with $d = 2g + 1$; we will adhere to this assumption about f throughout most of the work (see §4 for more details). Our aim will be showing how to explicitly form semistable models of Y over finite extensions of K. We more fully explain various aspects of the problem below.

1.1. Semistable models of curves. Given any smooth projective geometrically connected curve C over a complete discrete valuation field K with ring of integers $R \subset K$ and algebraically closed residue field k, a *model* of C over R', where R' is the ring of integers of some finite extension $K' \supseteq K$, is a normal projective flat R'-scheme $\mathcal{C}$ whose generic fiber is isomorphic to C over K'. We say that a model $\mathcal{C}$ is *semistable* if its special fiber $\mathcal{C}_s$ is a reduced k-curve with at worst nodes as singularities. The following groundbreaking theorem was proved by Deligne and Mumford in [6] and then through independent arguments by Artin and Winters in [1] (see also [11, Section 10.4] for a detailed explanation of the arguments in Artin-Winters).

THEOREM 1.1. Every smooth projective geometrically connected curve C over K achieves semistable reduction over a finite extension $K' \supseteq K$, i.e. C admits a semistable model $\mathcal{C}^{\mathrm{ss}}$ over R', where R' is the ring of integers in K'.

The above result is not constructive and does not tell us how to find a semistable model $\mathcal{C}^{\mathrm{ss}}$ or exactly how large an extension $K' \supseteq K$ is needed in order to define it. It moreover does not specify, for a given curve C/K, anything about the structure of the special fiber $(\mathcal{C}^{\mathrm{ss}})_s$. It is therefore natural to ask whether there is any general method by which we may construct a semistable model $\mathcal{Y}^{\mathrm{ss}}$ of a hyperelliptic curve Y/K defined by an equation of the form in (1).

1.2. Special fibers of semistable models of curves. In this work, we are interested not only in how to construct a semistable model $\mathcal{Y}^{\mathrm{ss}}$ of a hyperelliptic curve Y, but also in how certain characteristics of the defining polynomial may determine the *structure* of the special fiber of such a semistable model. The special fiber $(\mathcal{Y}^{\mathrm{ss}})_s$ of a semistable model $\mathcal{Y}^{\mathrm{ss}}$ of a curve Y/K by definition consists of reduced components which meet each other only at nodes. Each node, viewed as a point in $\mathcal{Y}^{\mathrm{ss}}$, has a *thickness* (see the initial discussion in §2.1.1.6) which is a positive integer. The structure of the special fiber $(\mathcal{Y}^{\mathrm{ss}})_s$ can be described entirely in terms of the set of its irreducible components, the genus of the normalization of each of these components, the data of which components intersect which others at how many nodes, and the thicknesses of the nodes. The sum of the genera of the normalizations of the irreducible components is known as the *abelian rank* of $(\mathcal{Y}^{\mathrm{ss}})_s$, while the number of loops in the configuration of components and their intersections (i.e. the number of loops in the dual graph of $(\mathcal{Y}^{\mathrm{ss}})_s$) is known as the *toric rank* of the special fiber of $\mathcal{Y}^{\mathrm{ss}}$. The property of being semistable implies that the sum of these two ranks equals the genus of Y. See §2.1.1.7 below for more details.

Replacing a semistable model $\mathcal{Y}^{\mathrm{ss}}$ of Y over R' with another semistable model of Y over R'' (where R' and R'' are the ring of integers of possibly different extensions of K) does not affect its abelian or toric rank (see Proposition 2.6 below), and therefore these ranks are intrinsic to the curve Y itself and particularly interesting to determine (meanwhile, the thicknesses of the nodes change in a predictable manner between semistable models over different extensions of R; see the discussions in §2.1.1.6 and §2.1.1.8).

1.3. The reduction of a curve given by $y^2 = f(x)$. Our first naïve attempt to produce a semistable model for Y is to perform simple changes of variables (if necessary) over a low-degree field extension $K' \supset K$ so that the coefficients appearing in the equation in (1) are all integral and then to simply use this equation to define a scheme $\mathcal{Y}$ over the corresponding ring of integers R'. More precisely, it is clear that after possibly scaling x and y by appropriate elements of $\bar{K}^\times$, we may assume that f is monic (i.e., $c = 1$), and that the roots a_i are all integral, with $\min_{i,j} v(a_i - a_j) = 0$. In particular, f has integral coefficients, and so we may extend Y to a scheme $\mathcal{Y}/R'$ whose generic fiber is Y and whose special fiber $\mathcal{Y}_s$ is given (over the affine chart $x \neq \infty$ of $\mathbb{P}^1_k$) by the equation

$$y^2 = \bar{f}(x) := \prod_{i=1}^{2g+1} (x - \bar{a}_i), \tag{2}$$

where each element $\bar{a}_i$ is the reduction of $a_i \in \mathcal{O}_{\bar{K}}$ in the residue field k.

Suppose that the residue characteristic of K is different from 2. Then the curve $\mathcal{Y}_s/k$ is generically an étale double cover of the projective line $\mathbb{P}^1_k$, and its only possible singularities are produced by multiple roots of the reduced polynomial $\bar{f}$; consequently, the reduced curve $\mathcal{Y}_s$ is smooth if and only if the roots of f are all distinct modulo the prime ideal of the splitting field.

Suppose on the other hand that the residue characteristic of K is 2. Then the curve $\mathcal{Y}_s/k$ is an inseparable cover of $\mathbb{P}^1_k$, and it *always* has non-nodal singularities whether or not the reduction of the polynomial f has multiple roots. We summarize these (fairly elementary) facts in the following proposition.

Proposition 1.2. Let Y/K and $\mathcal{Y}/R'$ be defined as in the discussion above.

(a) Suppose that the residue characteristic of K is not 2. Then each singular point of the special fiber $\mathcal{Y}_s$ is of the form $(x, y) = (\bar{a}, 0)$, where $\bar{a} \in k$ is a multiple root of the reduced polynomial $\bar{f}$. Given a singular point $(\bar{a}, 0)$ of $\mathcal{Y}_s$, let $\mathfrak{s} \subset \bar{K}$ be the subset of roots of f which each reduce to $\bar{a}$. Then,
   (i) if $\mathfrak{s}$ has cardinality 2, the singular point $(\bar{a}, 0)$ is a node; and
   (ii) if $\mathfrak{s}$ has cardinality at least 3, the singular point $(\bar{a}, 0)$ is not a node.

(b) Suppose that the residue characteristic of K is 2. Then the special fiber $\mathcal{Y}_s$ has a non-nodal singularity at each point whose x-coordinate is a root of the derivative polynomial $\bar{f}'$ (and these are the only singularities of $\mathcal{Y}_s$).

Proof. It is straightforward to verify, using a standard equation for another affine open subset of $\mathcal{Y}_s$ (which is given in §4) which contains the points over $x = \infty$, that there is no singular point over $x = \infty$ due to the fact that f is monic so that its reduction $\bar{f}$ has maximal degree. To prove both parts of the proposition, it therefore suffices to consider singular points on the affine part of $\mathcal{Y}_s$ defined by the equation $y^2 = \bar{f}(x)$.

Assume first that the residue characteristic is not 2. Then, applying the Jacobian criterion and setting both partial derivatives of $y^2 - \bar{f}(x)$ to 0, we get that a singular point can only occur where $y = 0$ (which implies that x is a root of $\bar{f}$) and x is a root of $\bar{f}'$. These conditions imply that the x-coordinate of a singular point must be a multiple root of $\bar{f}$. After appropriately translating the x-coordinate, we may assume that a given singular point is $(0, 0)$, which implies that $\bar{f}(x)$ is exactly divisible by x^n for some integer $n \geq 2$ which is the multiplicity of the root 0. The singular point $(0, 0)$ is a node if and only if the polynomial consisting of the terms of degree ≤ 2 in the defining polynomial $y^2 - \bar{f}(x)$ factors into distinct linear polynomials over k (this is indeed how *singular point* and *node* are defined in [15, §I.1.2]). This clearly happens if and only if $n = 2$, which finishes the proof of part (a).

Now assume that the residue characteristic is 2. This time, applying the Jacobian criterion tells us that there is a singular point wherever we have $\bar{f}'(x) = 0$. After translating both coordinate variables x and y suitably, we may assume that a given singular point is $(0, 0)$. Now it is clear that the polynomial consisting of the terms of degree ≤ 2 in the defining polynomial $y^2 - \bar{f}(x)$ does not factor into distinct linear polynomials over k since it does not include an xy-term and is therefore the square of a linear polynomial instead. This implies that the singular point is not a node, and part (b) is proved. $\square$

1.4. Cluster data. We have just seen that the naïve attempt to construct a semistable model of a hyperelliptic curve Y as in §1.1.1.3 always fails over residue characteristic 2. Meanwhile, in the case that the residue characteristic is not 2, Proposition 1.2 more or less implies that the naïve model $\mathcal{Y}/R'$ is semistable if and only if (1) the roots of f are *equidistant* (i.e. the valuations of the difference between the roots are all equal) so that $\mathcal{Y}_s$ is smooth, or (2) the roots of f are equidistant except for certain pairs of roots of f which are closer to each other with respect to the discrete valuation of K (so that each pair maps to a root of multiplicity 2 of the reduced polynomial $\bar{f}$ and produces a node of $\mathcal{Y}_s$). This suggests that when the residue characteristic of K is not 2, the data of the valuations of differences between roots of f may be directly crucial for constructing a semistable model of Y and for understanding the structure of the special fiber of such a semistable model.

This notion is made precise in [7] by defining the *cluster data* associated to a hyperelliptic curve Y over a discrete valuation field K: roughly speaking, if Y is defined by an equation of the form in (1), its associated cluster data consists of subsets $\mathfrak{s}$ of roots of the defining polynomial f, called *clusters*, which are closer to each other with respect to the discrete valuation of K than they are to the roots of f which are not contained in $\mathfrak{s}$, along with, for each non-singleton cluster $\mathfrak{s}$, the minimum valuation of differences between roots in $\mathfrak{s}$, called the *depth* of $\mathfrak{s}$. For precise definitions, see Definition 5.2 below or [7, Definition 1.1].

When the residue characteristic is different from 2, the process of construction of a semistable model of Y as well as the structure of its special fiber is governed entirely by the cluster data associated to Y. This can be deduced from the explicit constructions given in [7, §4, 5] in any case, but we will present a variant of this construction in §5. The rough idea is summarized as follows, under the simplifying assumption that f has degree $2g + 1$.

(1) There is a one-to-one correspondence between discs $D \subset \bar{K}$ (with respect to the induced valuation on $\bar{K}$) and smooth models of $\mathbb{P}^1_K$ over finite extensions of R, and each model of $\mathbb{P}^1_K$ over a finite extension of R with reduced special fiber is the compositum of a finite number of smooth models and thus corresponds to a finite collection of discs $D \subset \bar{K}$ (see §4.2 for more details).

(2) We define $\mathcal{X}^{(\mathrm{ss})}$ to be the model of $\mathbb{P}^1_K$ over a finite extension of R corresponding in the above way to the set of discs $D_{\mathfrak{s}}$ for all non-singleton clusters $\mathfrak{s}$, where each

$D_{\mathfrak{s}} \subset \bar{K}$ is defined to be the minimal disc whose intersection with the set of roots of f coincides with $\mathfrak{s}$.

(3) There is a semistable model $\mathcal{Y}^{\mathrm{ss}}$ with a degree-2 map to $\mathcal{X}^{(\mathrm{ss})}$ which is constructed simply by normalizing $\mathcal{X}^{(\mathrm{ss})}$ in the function field $K(Y)$ after possibly replacing K with a finite extension, which is at most the unique quadratic extension of the splitting field of f.

The discs D_i mentioned in Step (2) correspond to changes in coordinate of the form $x = \alpha_i + \beta_i x_i$ for some $\alpha_i \in \bar{K}$ and $\beta_i \in \bar{K}^\times$; for each such change in coordinate, we may perform appropriate substitutions into the equation $y^2 = f(x)$ and transform y appropriately to get a new equation of the form $y_i^2 = f_i(x_i) \in R'[x_i]$, where R' is the ring of integers of an appropriate finite extension $K' \supseteq K$; this new equation defines a model of Y, which is the normalization of the model of $\mathbb{P}^1_K$ corresponding to the disc D_i in the function field $K'(Y)$. The desired semistable model $\mathcal{Y}^{\mathrm{ss}}$ is comprised of these normalizations. The idea is illustrated by the following example.

Example 1.3. Let $K = \mathbb{Q}_p^{\mathrm{unr}}$ for some $p \geq 5$ and

$$f(x) = x(x-p^3)(x-p)(x-1)(x-1+p^4)(x-2)(x-3). \tag{3}$$

The set of roots of f is $\mathcal{R} := \{0, p^3, p, 1, 1-p^4, 2, 3\}$. The *clusters* of these roots (i.e. the subsets $\mathfrak{s}$ consisting of roots which are closer to each other than they are to the roots in $\mathcal{R} \smallsetminus \mathfrak{s}$) are

$$\mathfrak{s}_0 := \mathcal{R},\ \mathfrak{s}_1 := \{0, p^3, p\},\ \mathfrak{s}_2 := \{0, p^3\},\ \mathfrak{s}_3 := \{1, 1-p^4\},$$

as well as each of the singleton subsets of $\mathcal{R}$ (which we ignore). The data of these clusters is represented by the following diagram.

$$\text{cluster picture of } \mathcal{R}: \quad \boxed{\boxed{\boxed{0 \qquad p^3}_{\mathfrak{s}_2} \qquad p}_{\mathfrak{s}_1} \qquad \boxed{1 \qquad 1-p^4}_{\mathfrak{s}_3} \quad 2 \qquad 3}_{\mathfrak{s}_0}$$

The discs $D_i \subset \bar{K}$ minimally containing each of the clusters $\mathfrak{s}_i$ are then given by

$$\begin{aligned} &D_0 := \overline{\mathbb{Z}_p}, && D_1 := p\overline{\mathbb{Z}_p} = \{0 + pz \mid z \in \overline{\mathbb{Z}_p}\}, \\ &D_2 := p^3\overline{\mathbb{Z}_p} = \{0 + p^3 z \mid z \in \overline{\mathbb{Z}_p}\}, && D_3 := 1 + p^4\overline{\mathbb{Z}_p} = \{1 + p^4 z \mid z \in \overline{\mathbb{Z}_p}\}, \end{aligned} \tag{4}$$

where $\overline{\mathbb{Z}_p}$ denotes the ring of integers of the algebraic closure $\overline{\mathbb{Q}_p}$ of $\mathbb{Q}_p$. The changes in coordinates corresponding to each of these discs are given by

$$x = x_0 = px_1 = p^3 x_2 = p^4(x_3 - 1),$$

where each x_i corresponds to the disc D_i in an obvious way, and we define corresponding coordinates y_i by scaling y by suitable elements of $\mathbb{Q}_p(\sqrt{p})$ as

$$y = y_0 = p^{3/2} y_1 = p^{7/2} y_2 = p^4 y_3.$$

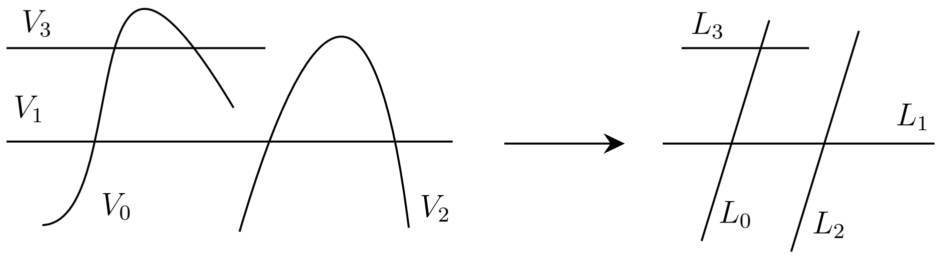

FIGURE 1. The special fiber $(\mathcal{Y}^{\mathrm{ss}})_s$, shown on the left, mapping to the special fiber $(\mathcal{X}^{\mathrm{ss}})_s$; each component V_i of $(\mathcal{Y}^{\mathrm{ss}})_s$ maps to each component $L_i := (\mathcal{X}_{D_i})_s$ of $(\mathcal{X}^{\mathrm{ss}})_s$.

We now define corresponding models $\mathcal{Y}_i/\mathbb{Z}_p^{\mathrm{unr}}[\sqrt{p}]$ of $Y/\mathbb{Q}_p^{\mathrm{unr}}(\sqrt{p})$ for $i = 0, 1, 2, 3$, given by the below equations.

(5)

$$\begin{aligned}
&\mathcal{Y}_0 : y_0^2 = f(x) = f(x_0)\\
&\mathcal{Y}_1 : y_1^2 = p^{-3}f(x) = x_1(x_1 - p^2)(x_1 - 1)(px_1 - 1)(px_1 - 1 + p^4)(px_1 - 2)(px_1 - 3)\\
&\mathcal{Y}_2 : y_2^2 = p^{-7}f(x) = x_2(x_2 - 1)(p^2x_2 - 1)(p^3x_2 - 1)(p^3x_2 - 1 + p^4)(p^3x_2 - 1)(p^3x_2 - 2)\\
&\mathcal{Y}_3 : y_3^2 = p^{-8}f(x) = (p^4x_3 - 1)(p^4x_3 - 1 - p^3)(p^4x_3 - 1 - p)(x_3)(x_3 - 1)(p^4x_3 - 2)
\end{aligned}$$

Their respective reductions (that is, their special fibers $(\mathcal{Y}_i)_s$) over the residue field $\overline{\mathbb{F}_p}$ are as follows.

$$\begin{aligned}
&(\mathcal{Y}_0)_s : y_0^2 = x_0^3(x_0 - 1)^2(x_0 - 2)(x_0 - 3)\\
&(\mathcal{Y}_1)_s : y_1^2 = 6x_1^2(x_1 - 1)\\
&(\mathcal{Y}_2)_s : y_2^2 = -6x_2(x_2 - 1)\\
&(\mathcal{Y}_3)_s : y_3^2 = 2x_3(x_3 - 1)
\end{aligned} \tag{6}$$

The desingularizations of each of these special fibers give rise to the components of the special fiber of the desired semistable model $\mathcal{Y}^{\mathrm{ss}}$: here $(\mathcal{Y}_0)_s$ contributes a smooth component V_0 of genus 1; $(\mathcal{Y}_1)_s$ contributes a line V_1 which intersects V_0 at a single node; $(\mathcal{Y}_2)_s$ contributes a line V_2 which intersects V_1 at 2 nodes; and $(\mathcal{Y}_3)_s$ contributes a line V_3 which intersects V_0 at 2 nodes. The configuration is shown in Figure 1.

One can see from the configuration of components displayed in Figure 1 that the toric rank of $(\mathcal{Y}^{\mathrm{ss}})_s$ is 2; if one adds this to the sum of the genera of the components V_i, the genus $g = 3 = 2 + 1$ of Y is recovered.

Remark 1.4. In the case that Y/K is an *elliptic curve* (i.e. $g = 1$) over residue characteristic $p \neq 2$, where the polynomial f has degree 3, there are at most 2 non-singleton clusters of roots of f, and a similar procedure can be performed to get a semistable model of Y over the (unique) quadratic ramified extension of the splitting field of f, which will be smooth if and only if Y/K has potentially good reduction. This is more or less the process outlined in the proof of [16, III.1.7(a)] combined with the proof of [16, VII.5.4(c)], except that Silverman does not construct a separate component of the semistable model corresponding to a cardinality-2 cluster of roots (in the case that there is one). So, following Silveman's method, the special fiber of the semistable model always consists of only 1 component which has a node if and only if there is a cardinality-2 cluster of roots of f (this is the case of *multiplicative reduction*).

When the residue characteristic of K is 2, it is natural to ask whether a semistable model of Y can be constructed by a procedure governed entirely by the associated cluster data in this way. In short, the answer is "no", but in this work we develop methods of finding a particular collection of discs in $\bar{K}$ which corresponds to a model $\mathcal{X}^{(\mathrm{ss})}$ of $\mathbb{P}^1_K$ over a finite extension of R, such that the model $\mathcal{Y}^{\mathrm{ss}}$ of Y which is constructed directly from $\mathcal{X}^{(\mathrm{ss})}$ in a similar manner to Steps (2)-(3) above is guaranteed to be semistable (and to satisfy several other nice properties discussed in §3). We will present and prove results relating such a set of discs to the set of clusters $\mathfrak{s}$ appearing in the cluster data associated to Y.

2. A summary of our main results for residue characteristic 2

Although the arguments used in this work will recover what is already known about the construction of semistable models of hyperelliptic curves in characteristic different from 2, our primary aim is to understand how to construct a semistable model as well as the structure of its special fiber when the residue characteristic is 2. This is addressed by our main results.

2.1. Constructing equations for models with semistable reduction. It is clear from Proposition 1.2 that if K has residue characteristic 2, a model of Y given by an equation of the form $y^2 = f(x) \in R[x]$ cannot possibly have semistable reduction. We must therefore find a model given by one or more equations of the more general form

$$y_i^2 + q_i(x_i)y_i = r_i(x_i), \tag{7}$$

where $q_i(x_i), r_i(x_i) \in R'[x_i]$ are polynomials of degree less than or equal to $g+1$ and $2g+1$ respectively (see §4.1 below for more details on this form of equation). This is generally accomplished in the following manner. First (as in the case of residue characteristic not 2) we make a substitution of the form $x = \alpha_i + \beta_i x_i$ with $\alpha_i \in \bar{K}$ and $\beta_i \in \bar{K}^\times$ and scale y by a suitable element of $\bar{K}^\times$ to get a coordinate $\tilde{y}_i$ and an equation of the form $\tilde{y}_i^2 = f_i(x_i) \in \bar{K}[x_i]$, where f_i has integral coefficients and nonzero reduction. Then, in order to turn this into an equation of the form in (7), we find a decomposition $f_i = q_i^2 + 4r_i$, where $q_i(x_i), r_i(x_i) \in \bar{K}[x_i]$ are polynomials of degree less than or equal to $g+1$ and $2g+1$ respectively (this is a *part-square decomposition* as we define it below in Definition 4.13) and set $\tilde{y}_i = 2y_i + q_i(x)$; note that this is essentially performing the standard operation of "completing the square" in reverse.

There are two points of delicacy that must be taken into account when choosing the elements α_i, β_i and the decomposition $f_i = q_i^2 + 4r_i$. One is that α_i and β_i must be chosen carefully so that all terms in the resulting equation of the form (7) have integral coefficients, so that these equations may be defined over the ring of integers $R' \supseteq R$ of the finite extension of K given by adjoining all elements α_i, β_i and coefficients of the polynomials q_i, r_i. Secondly, one must be sure that all of the components of $(\mathcal{X}^{\mathrm{ss}})_s$ (each corresponding to a choice of α_i and β_i) have really been found; otherwise, the model of Y

corresponding to an incomplete set of coordinates x_i will contain non-nodal singularities in its special fiber and will therefore fail to be semistable.

2.2. Our main results. As discussed above, a semistable model $\mathcal{Y}^{\mathrm{ss}}/R'$ of Y (where R' is the ring of integers of a finite extension K'/K) may be constructed, more or less, as the normalization of a suitable model $\mathcal{X}^{(\mathrm{ss})}$ of $\mathbb{P}^1_K$ in the function field of Y, and $\mathcal{X}^{(\mathrm{ss})}$, in turn, corresponds to a finite collection of changes of coordinate $x = \beta_i x_i + \alpha_i$ (so that its special fiber is composed of copies of the projective line over the residue field k corresponding to each coordinate x_i; see §4.2 below). Finding a collection of appropriate substitutions $x = \beta_i x_i + \alpha_i$ is therefore in some sense the most essential step in finding our desired semistable model $\mathcal{Y}^{\mathrm{ss}}$, just as it is in the case of residue characteristic not 2. As in our discussion in §1.1.1.4, each new coordinate x_i obtained from x in this way by translation and homothety corresponds to a disc $D_i := \{\beta_i z + \alpha_i \mid z \in \mathcal{O}_{\bar{K}}\}$, so finding a semistable model of Y again largely amounts to choosing an appropriate collection of discs in $\bar{K}$. The difference now is that, unlike in the case of residue characteristic not 2, these discs generally do not correspond in a one-to-one manner to non-singleton clusters of roots of f.

In §3.4 of this work, we define a particularly nice (unique up to unique isomorphism) semistable model $\mathcal{Y}^{\mathrm{rst}}$ of a given hyperelliptic curve Y which we call the *relatively stable model* (see Definition 3.8 below). We will define a *valid disc* (Definition 5.11 below) to be a disc $D \subset \bar{K}$ among the collection of discs used the manner discussed above to construct the semistable model $\mathcal{Y}^{\mathrm{rst}}$ (excluding such discs which correspond to components of $(\mathcal{X}^{(\mathrm{rst})})_s$ over which the cover $(\mathcal{Y}^{\mathrm{rst}})_s \to (\mathcal{X}^{(\mathrm{rst})})_s$ is inseparable). The central results we present in this work are on how to find valid discs. While the exact procedure provided by these results cannot be described succinctly in this introduction, we give a partial summary of the general outcome in the following theorem.

THEOREM 1.5. Assume all of the above set-up for a hyperelliptic curve Y/K of genus g given by an equation of the form $y^2 = f(x) \in K[x]$, where the polynomial f has degree $2g+1$, and assume that the residue characteristic of K is 2. Let $\mathcal{Y}^{\mathrm{rst}}/R'$ be the relatively stable model of Y, where R' is the ring of integers of an appropriate finite field extension $K' \supseteq K$. Let $\mathcal{R} \subset \bar{K}$ denote the set of roots of f. For any cluster of roots $\mathfrak{s} \subsetneq \mathcal{R}$, we write $\mathfrak{s}'$ for the minimal cluster which properly contains $\mathfrak{s}$.

The clusters of roots in $\mathcal{R}$ and the valid discs associated to Y are related in the following manner.

(a) Given a valid disc $D \subseteq \bar{K}$, the cardinality of $D \cap \mathcal{R}$ is even (and we may have $D \cap \mathcal{R} = \varnothing$).

(b) If a cluster $\mathfrak{s}$ has even cardinality, there are either 0, 1, or 2 valid discs $D \subseteq R'$ such that either $D \cap \mathcal{R} = \mathfrak{s}$ or D is the smallest disc containing $\mathfrak{s}'$.

(c) Let $\mathfrak{s}$ be an even-cardinality cluster of *relative depth* $m := \min\{v(a-a') \mid a, a' \in \mathfrak{s}\} - \min\{v(a-a') \mid a, a' \in \mathfrak{s}'\}$ (see Definition 5.1), and write $f_0(x) = \prod_{a \in \mathfrak{s}}(x-a)$ and $f_\infty(x) = f(x)/f_0(x)$. There exists a rational number $B_{f,\mathfrak{s}} \in \mathbb{Q}_{\geq 0}$ which is independent of the relative depth of $\mathfrak{s}$ in the sense of Remark 6.27, such that
   (i) if $m > B_{f,\mathfrak{s}}$, the number of valid discs as in part (b) is "2";
   (ii) if $m = B_{f,\mathfrak{s}}$, the number of valid discs as in part (b) is "1"; and
   (iii) if $m < B_{f,\mathfrak{s}}$, the number of valid discs as in part (b) is "0".

Moreover, in the case of (i), the 2 guaranteed valid discs containing $\mathfrak{s}$ give rise to 2 components of $(\mathcal{Y}^{\mathrm{rst}})_s$ which intersect each other at 2 nodes, each having thickness equal to $(m - B_{f,\mathfrak{s}})/v(\pi)$, where π is a uniformizer of K'.

(d) Given an even-cardinality cluster $\mathfrak{s}$, the bound $B_{f,\mathfrak{s}}$ from part (d) satisfies $B_{f,\mathfrak{s}} \leq 4v(2)$. If we furthermore assume that $\mathfrak{s}$ and $\mathfrak{s}'$ each have a maximal subcluster of odd cardinality (e.g. a maximal subcluster which is a singleton), we have the inequality

$$B_{f,\mathfrak{s}} \geq \Big(\frac{2}{|\mathfrak{s}|-1} + \frac{2}{2g+1-|\mathfrak{s}|}\Big)v(2). \tag{8}$$

(e) The toric rank of some (any) semistable model of Y is equal to the number of even-cardinality clusters satisfying item (i) above which themselves cannot be written as a disjoint union of such even-cardinality clusters.

(f) Let $\mathfrak{s}$ be a cluster of odd cardinality not equal to 1 or $2g+1$. Then $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of two curves C_0 and C_∞ meeting as a single node in $(\mathcal{Y}^{\mathrm{rst}})_s$; their arithmetic genera are $\frac{1}{2}(|\mathfrak{s}|-1)$ and $g - \frac{1}{2}(|\mathfrak{s}|-1)$ respectively.

The statements in the above theorem are a combination of a (sometimes simplified version of) statements of the main results presented and proved in this work. Parts (a)–(c), apart from the final statement in (c), are adapted from Theorem 6.18 and Proposition 6.26 below (see also Theorem 5.13); part (d) is adapted from Proposition 6.35(c); the final statement in (c) comes directly from Proposition 8.4; part (e) is a rephrasing of Theorem 8.1; and part (f) is a rephrasing of Corollary 8.19 (we note that this statement actually also holds when the residue characteristic is different from 2). Formulas for thicknesses are not explicitly given in the above-mentioned results but in general can easily be computed using Proposition 3.4(b) combined with Proposition 4.8; we get the assertion about thicknesses in part (c) from applying Proposition 8.5 to our results in §6.3 (see Proposition 6.24) which tell us explicitly what the depths of the 2 guaranteed valid discs are in the situation of Theorem 1.5(c)(i).

The results in this work can be viewed as a vast generalization of the results in [17], where J. Yelton explicitly constructed semistable models of elliptic curves with a cluster of cardinality 2 and depth m (as well as elliptic curves with no even-cardinality clusters). The threshold for m above which there are 1 or 2 valid discs containing that cardinality-2 cluster which is found in [17] comes as the following easy corollary to the above theorem; we remark that this corollary can be deduced also from standard formulas for the j-invariant of an elliptic curve (specifically, the particular choice of power of 2 multiplied to the rest of the formula, which influences the valuation of the j-invariant in residue characteristic 2; see Remark 9.5(a) below).

Corollary 1.6. Suppose that we are in the $g = 1$ case of the situation in Theorem 1.5 and that $\mathfrak{s}$ is a cluster of cardinality 2. Then we have $B_{f,\mathfrak{s}} = 4v(2)$.

Proof. The parent cluster of $\mathfrak{s}$ (i.e., the minimal cluster strictly containing) is $\mathfrak{s}' = \mathcal{R}$, which has cardinality 3. It is clear that both $\mathfrak{s}$ and $\mathfrak{s}'$ have a singleton child cluster (i.e., a maximal subcluster consisting of only one root). Now, Theorem 1.5(d) gives that $B_{f,\mathfrak{s}} \leq 4v(2)$ and

$$B_{f,\mathfrak{s}} \geq \big(\frac{2}{1} + \frac{2}{1}\big)v(2) = 4v(2). \tag{9}$$

The equality $B_{f,\mathfrak{s}} = 4v(2)$ follows. $\square$

For examples of semistable models of hyperelliptic curves over residue characteristic 2 which are explicitly computed in the manner discussed above, see Examples 9.6 and 9.13 below, which are worked out directly from the results and processes developed in §6 and §7. Note that in both of these examples, the set of clusters consists of a single cardinality-2 cluster $\mathfrak{s}$ as well as the full set $\mathcal{R}$ of the roots, so that following what happens in the case of residue characteristic not 2, we would expect that $(\mathcal{Y}^{\mathrm{rst}})_s$ contains exactly 2 components obtained by centering at an element of the $\mathfrak{s}$ and scaling according to how close the 2 elements in $\mathfrak{s}$ are. However, in this case, the choices of scaling factors β_i are not so "obvious" as in the situation of residue characteristic not 2 (as in Example 1.3), and moreover, in Example 9.13 we get a further component.

Theorem 1.5 above describes the overall relationship between clusters and valid discs associated to a hyperelliptic curve over residue characteristic 2, which is one of our main points of focus, but in our more broad investigation we come up with a general method of finding all valid discs. The process of finding all valid discs having a given center (in particular, those containing a given cluster) is developed in §6 (the actual computations that are necessary are aided by Algorithm 6.43), while for residue characteristic 2, the process of finding centers of all valid discs (in particular the ones which do not contain roots of f) is developed in Chapter 7, relying on the computation of a certain polynomial $F(T) \in K[T]$; Corollary 7.8(a) states in particular that each valid disc not containing roots of f is centered at a root of F.

3. Comparison to other works

A hyperelliptic curve is a special case of a *superelliptic curve*, i.e. a curve defined by an equation of the form $y^n = f(x)$ for some $n \geq 2$. There have been a number of works discussing semistable models of superelliptic curves. When the exponent n in the equation for a superelliptic curve is not divisible by the residue characteristic p, the process of constructing a semistable model is relatively straightforward and is provided in [4, §3], [3, §4], [7, §4, 5] (for hyperelliptic curves, using the language of clusters), and [9] (for hyperelliptic curves, using the language of stable marked curves), as well as earlier works. We recover our own variant of their results in the hyperelliptic case (i.e. when $n = 2$ and $p \neq 2$) based on Theorem 5.12 below, in the process of investigating the situation when $p = 2$.

The existing results for the wild case of semistable reduction of superelliptic curves, i.e. when the defining equation is of the form $y^p = f(x)$ where p is the residue characteristic, have been far more limited. To the best of our knowledge, investigations into this case began with Coleman, who in [5] outlined an algorithm for changing coordinates in such a way that the defining equation is converted to a form whose reduction over the residue field does not describe a curve which is an inseparable degree-p cover of the line; when $p = 2$, this is more or less equivalent to our notion of *part-square decompositions* which will be introduced in §4.3. This idea is further developed by Lehr and Matignon in [13] and later in [10] (among several other works). Their results apply only to the very particular case of *equidistant geometry*, meaning that the valuations of differences between each pair of distinct roots of the defining polynomial f are all equal, which in the language of clusters means that there are no proper, non-singleton clusters of roots. Much of their focus is on the (finite) extension of the ground field over which semistable reduction is obtained and the action of the (finite) Galois group of this extension on the special fiber of their semistable model. The wild case is also discussed in [3, §4], in which several examples are computed and interpreted in terms of rigid analytic geometry; the working of these examples is mainly done through clever guessing rather than a direct algorithm, however.

There are further similarities between the ideas presented in the work of Lehr and Matignon and some of our results, which are applied to hyperelliptic curves whose branch points are not necessarily geometrically equidistant. The notion of *p-développements de Taylor* (Taylor p-expansions) introduced in [13, §2], while defined completely differently, is alike in motivation and applications to our notion of *sufficiently odd decompositions* (see §6.5 below), and our algorithm for computing sufficiently odd decompositions is a mild variation of [13, Proposition 2.2.1], which is used to show that *p-développements de Taylor* exist. Moreover, in each of [13] and [10], a polynomial over the ground field is defined whose roots are the centers of all discs which give rise to components of the special fiber; these polynomials (the *p-dérivée* in [13, Définition 2.4.1] and the *monodromy polynomial* in [10, Definition 3.4]) are quite distinct but each is defined similarly and plays a similar role to our polynomial $F(T) \in K[T]$ given in Definition 7.2 below, whose roots in the geometrically equidistant case certainly provide centers of all the valid discs.

Our work differs from the prior research discussed above in that our major focus is on the relationship between clusters of roots and the structure of the special fiber of a semistable model of a hyperelliptic curve when the residue characteristic is 2; to the best of our knowledge, the only specific case in terms of cluster data which has been investigated where equidistant geometry is not assumed is in the recent preprint [8], which treats a case involving an even number of roots clustering in pairs.

We finish this subsection by remarking that our work does not prioritize much focus towards describing the finite extension of K over which we are constructing our semistable (relatively stable) model of Y or determining the minimal extension of K over which Y achieves semistable reduction (although in building $\mathcal{Y}^{\mathrm{rst}}$, we try to be economical in the extension of K required). However, as our results are constructive, it is fairly straightforward to compute the (necessarily totally ramified) extension K'/K over which $\mathcal{Y}^{\mathrm{rst}}$ is defined. In general, the extension K'/K is obtained from (possibly) a sequence of quadratic extensions of the subfield $K'' \subset K'$ over which the associated model $\mathcal{X}^{(\mathrm{rst})}$ of the projective line is defined using changes of coordinates $x = \alpha_i + \beta_i x_i$ as discussed above; then K'' is clearly just the smallest field over which the discs of $\mathfrak{D}^{(\mathrm{rst})}$ are defined, where a disc D of $\bar{K}$ is said to be *defined* over a field extention K'/K if there exist α and β in K' such that $D = D_{\alpha,v(\beta)}$. In practice, each scaling element β_i may be chosen to be any element of a prescribed valuation, while a given translating element α_i may be chosen to be a root of f (and thus already in the splitting field) when the corresponding valid disc contains roots of f; it is only in the case where there are valid discs not containing roots of f that one may have to choose α_i to be a root of the (generally high-degree) polynomial $F(T) \in K[T]$ defined in §7.1. It would be interesting to pursue results that specify the minimal extension K'/K over which $\mathcal{Y}^{\mathrm{rst}}$ (or some semistable model of Y) is defined under various hypotheses on f (or specify only its degree or its maximal tame subextension) and apply such results to other arithmetic questions (for instance, involving division fields of the Jacobian variety of Y).

4. Outline

While our priority in this work is considering the case where the residue characteristic p of our ground field K is 2, we try to be as general as possible so that we may at times compare and contrast the situation of $p = 2$ with the situation of $p \neq 2$, often considering the latter as a special case which yields more simply-stated results. We shall state the results for $p \neq 2$ using our own set-up and terminology which arises naturally from our method of recovering them but, as we have already mentioned, equivalent results do already appear in the literature (particularly in [7]). Beginning in §4 and throughout the rest of the work, we often refer to the two cases as "the $p = 2$ setting" and "the $p \neq 2$ setting".

The rest of our work is organized as follows. First, we establishing the algebro-geometric setting that we need in §2, which begins with briefly providing the basic background definitions and facts relating to models of curves over local rings, and then proceeds to look more closely at the properties of the special fiber of such a model and how to compare two models of the same curve by considering (-1)-lines and (-2)-curves (see Definitions 2.3 and 2.9 below). All of this set-up allows us in the following section to define a particular "nice" semistable model of a curve Y which is a Galois cover of another curve X, which we call the *relatively stable model* $\mathcal{Y}^{\mathrm{rst}}$ of Y (see Definition 3.8 below) and which is the main topic of §3. Viewing a hyperelliptic curve Y/K as a degree-2 (Galois) cover of the projective line $\mathbb{P}^1_K =: X$, the relatively stable model of Y is the one directly treated in the rest of this work.

After this rather general set-up, we specialize to considering models of hyperelliptic curves over discrete valuation rings. As a hyperelliptic curve is (by definition) a double cover of a projective line, we first look at models of projective lines over discrete valuation rings; the well-known characterization of such models is summarized in §4.1. Then in the rest of §4, we look at models of hyperelliptic curves from the point of view of algebraic equations which define them. More precisely, we derive equations which define normalizations of smooth models of the projective line (possibly looking over finite extensions of K) in the function field of the hyperelliptic curve Y. In the $p = 2$ setting, we describe how we use *part-square decompositions* (see Definition 4.13 below) of the defining polynomial of Y to find these normalizations.

We next turn our attention to clusters in §5, laying out the definitions of *clusters* and *cluster data* as in [7] (and other subsequent works) as well as introducing *valid discs* (see Definition 5.11 below), which by definition correspond more or less to the smooth models of the projective line comprising the model $\mathcal{X}^{(\mathrm{rst})}$ of the projective line of which the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ is the normalization in $K(Y)$. In this section, we essentially recover (as Theorem 5.12) the method of constructing a semistable model of Y according to cluster data in the $p \neq 2$ setting by showing that in this case, there is a one-to-one correspondence between valid discs and non-singleton clusters. The closest that we can come to an analogous statement for the $p = 2$ setting is then presented as Theorem 5.13 (which provides some of the statements of Theorem 1.5), but we defer the proof this theorem to §6.

The next two sections of our work focus on developing a method of finding valid discs for any particular hyperelliptic curve Y. The objective of §6 is an investigation of how to determine the existence and find the radius of a valid disc with a given center, whereas the goal in §7 is to show how to find those elements of $\bar{K}$ which are centers of valid discs. The main focus in §6 is on finding valid discs containing a given cluster of roots. In the course of developing the methods presented in this section, given a polynomial f, we define lower-degree polynomials $f_{+}^{\mathfrak{s}}$ and $f_{-}^{\mathfrak{s}}$ determined by a particular even-cardinality cluster $\mathfrak{s}$ of roots of f such that part-square decompositions of $f_{\pm}^{\mathfrak{s}}$ can be used to determine the existence and depths of valid discs containing $\mathfrak{s}$. One of the main findings is that an even-cardinality cluster $\mathfrak{s}$ has 2 associated valid discs if and only if the *depth* of $\mathfrak{s}$ exceeds a certain "threshold" $B_{f,\mathfrak{s}} \in \mathbb{Q}$ as in Theorem 1.5(c). In §6.4, we present and prove a number of results which give exact formulas or estimates of $B_{f,\mathfrak{s}}$ that apply to various situations, in particular proving the inequalities in Theorem 1.5(d). In §6.6 we present an algorithm for finding useful part-square decompositions of $f_{\pm}^{\mathfrak{s}}$ (Algorithm 6.43). Meanwhile, in §7, we characterize the centers of valid discs by defining a polynomial (Definition 7.2 below) whose roots are centers of all valid discs, with certain exceptions, as described by Theorem 7.6. The results of §6 and §7 together show how all valid discs may be found; in particular, Theorem 7.6 is guaranteed (by Corollary 7.8) to find centers of all valid discs which do not contain any roots of the defining polynomial f.

In §8, we proceed to examine the structure of the special fiber of our desired semistable model $\mathcal{Y}^{\mathrm{rst}}$ given knowledge of the valid discs containing particular clusters of roots. In this section, we show that in the situation of Theorem 1.5(c)(i) above, the guaranteed pair of valid discs, under certain circumstances, produces a loop in the graph of components of the special fiber of $\mathcal{Y}^{\mathrm{rst}}$, or in other words, increases the toric rank of the hyperelliptic curve by 1. This allows us to present (as Theorem 8.1) and prove a formula for the toric rank in terms of viable valid discs, as seen in Theorem 1.5(d).

Finally, we devote §9 to providing more direct formulas for the aforementioned polynomial F as well as the bounds $B_{f,\mathfrak{s}}$ for low-genus hyperelliptic curves, classified according to their associated cluster data (for the special case of genus 1, that is, for elliptic curves, this recovers the results which were presented and proved in a more concretely elementary way in [17]). In particular, Theorem 9.8 describes the possible structures of the special fiber of $\mathcal{Y}^{\mathrm{rst}}$ for genus-2 hyperelliptic curves classified according to their cluster data and broken into cases depending on valuations of certain elements of $\bar{K}$ associated to the defining polynomial.

5. Notations and conventions

Throughout this work, we adhere to the following notation:

- K is a field endowed with a discrete valuation $v : K \to \mathbb{Q} \cup \{+\infty\}$, complete with respect to v; we assume the residue field k of K is algebraically closed; when studying hyperelliptic curves over K (i.e., from §4 on), we will also always assume that $\mathrm{char}(K) \neq 2$;
- $R = \mathcal{O}_K = \{z \in K \mid v(z) \geq 0\}$ is the ring of integers of K;
- p is the characteristic of k (that is, p is the residue characteristic of K);
- thanks to the completeness of K, given any algebraic extension $K' \supseteq K$, the valuation $v : K \to \mathbb{Q} \cup \{+\infty\}$ extends uniquely to a valuation on K' which we also denote by $v : K' \to \mathbb{Q} \cup \{+\infty\}$: this turns K' into a non-archimedean field with residue field k, whose ring of integers will be denoted $R' := \mathcal{O}_{K'}$; when the

extension K'/K is finite, K' is actually a complete discretely-valued field, and R' is hence a complete DVR; and
- $\bar{K}$ is an algebraic closure of K.

Whenever we use interval notation (e.g. $[a, b]$, (a, b), $(a, +\infty)$, etc.), the bounds will always be elements of $\mathbb{Q} \cup \{\pm\infty\}$, and the interval will be understood to consist of all *rational* numbers (rather than all real numbers) between the bounds; i.e. we have $[a, b] = [a, b] \cap \mathbb{Q}$, $[a, +\infty] = [a, +\infty) = [a, +\infty) \cap \mathbb{Q}$, etc.

Let $h \in \bar{K}[z]$ be a polynomial; we denote its degree by $\deg(h)$. Then we again write $v : \bar{K}[z] \to \mathbb{Q} \cup \{+\infty\}$ for the *Gauss valuation*; that is, for any polynomial $h(z) := \sum_{i=0}^{\deg(h)} H_i z^i \in \bar{K}[z]$, we set

$$v(h) = v\left(\sum_{i=0}^{\deg(h)} H_i z^i\right) = \min_{1 \leq i \leq \deg(h)} \{v(H_i)\}.$$

5.1. Models. Given a smooth projective geometrically connected curve C over K, we typically use the same letter in curly font to denote a model $\mathcal{C}/R'$ defined over the ring of integers R' of some finite extension K'/K, and we denote its special fiber by $\mathcal{C}_s/k$. We denote by $\mathrm{Irr}(\mathcal{C}_s)$ the set of irreducible components of $\mathcal{C}_s$; to any given $V \in \mathrm{Irr}(\mathcal{C}_s)$, the notations $a(V), m(V)$ and $w(V)$ are defined in §2.2 (for the definition of the invariant $\underline{w}(V)$ in the Galois cover setting, see §3.2). Moreover, we denote by $\mathrm{Sing}(\mathcal{C}_s)$ the set of all singular points of the k-curve $\mathcal{C}_s$. We denote by $a(\mathcal{C}_s), t(\mathcal{C}_s), u(\mathcal{C}_s)$ the abelian, toric and unipotent rank of $\mathcal{C}_s$ (see Subsection 2.1.1.7). If $\mathcal{C}'$ is another model of $\mathcal{C}$, the notation $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}')$ denotes the set of points of $\mathcal{C}_s$ to which the irreducible components of $\mathcal{C}'_s$ that do not appear in $\mathcal{C}_s$ are contracted: see §2.1.1.3.

5.2. Lines and hyperelliptic curves. Beginning in §4, the symbol X will normally denote the projective line $\mathbb{P}^1_K$, and x will be its standard coordinate. Similarly, beginning in §4, the symbol Y will in general be used to denote a hyperelliptic curve of any genus $g \geq 1$ over K and ramified over $\infty \in X(K)$ and endowed with a 2-to-1 ramified cover map $Y \to X$; over the affine chart $x \neq \infty$, Y can be described by an equation of the form $y^2 = f(x)$, with $f(x) \in K[x]$ a polynomial of odd degree $2g + 1$. The set of the $2g + 1$ roots of $f(x)$ will be denoted $\mathcal{R} \subseteq \bar{K}$. We will use the notation $\mathcal{R} \cup \{\infty\}$ to mean the set of *all* $2g + 2$ branch points of $Y \to X$, including ∞.

In §3, we work with Galois covers in greater generality, and in that section $Y \to X$ indicates any Galois cover of smooth projective geometrically connected K-curves.

5.3. Discs and depths. In this work, we will often speak of the *depths* of clusters and of discs in $\bar{K}$; in the case of clusters, our definition of *depth* is the one used throughout [7]. More general, we can define depth for a subset of elements of $\bar{K}$ to measure how close the elements in the subset are to each other, as follows.

Definition 1.7. Given a subset $S \subset \bar{K}$, if a minimum of the valuations $v(\zeta - \zeta') \in \mathbb{Q} \cup \{+\infty\}$ among all elements $\zeta, \zeta' \in S$ exists, we call it the *depth* of S.

In this work, all depths will be rational numbers so that there will always exist an element of $\bar{K}$ whose valuation is equal to any given depth.

Given $\alpha \in \bar{K}$ and $\beta \in \bar{K}^{\times}$, we denote by $D_{\alpha,b}$ the subset of $\bar{K}$ described as $D_{\alpha,b} := \{x \in \bar{K} : v(x - \alpha) \geq b\}$: it can be characterized as the maximal subset of $\bar{K}$ containing α and having depth b. By a *disc* (of $\bar{K}$), we mean any subset of $\bar{K}$ of the form $D_{\alpha,b}$ for some $\alpha \in \bar{K}$ and $b \in \mathbb{Q}$. At various times when it is convenient, we will instead write $D_{\mathfrak{s},b}$, where $\mathfrak{s} \subset \bar{K}$ is a finite subset (in our contexts, it will often be a cluster of roots) rather than a single element and $b \in \mathbb{Q}$ is at least the depth of $\mathfrak{s}$, to denote the (necessarily unique) disc containing the subset $\mathfrak{s}$ and having depth b. Note that the depth of a disc is essentially minus a logarithm of its radius under the p-adic metric, and so a *greater* depth corresponds to a *smaller* disc.

Given a disc D of $\bar{K}$, the notation $\mathcal{X}_D$ denotes the corresponding smooth model of the line $X = \mathbb{P}^1_K$; given $\mathfrak{D}$ a finite non-empty collection of discs, we denote by $\mathcal{X}_{\mathfrak{D}}$ the corresponding model of the line $X = \mathbb{P}^1_K$ having reduced special fiber (see §4.2).

5.4. Translations and scalings. Given a variable z, an element $\alpha \in \bar{K}$ and an element $\beta \in \bar{K}^{\times}$, the notation $z_{\alpha,\beta}$ will be used to signify the translated and scaled variable $z_{\alpha,\beta} = \beta^{-1}(z - \alpha)$. Given a polynomial $h(z)$, we will denote by $h_{\alpha,\beta}$ the polynomial such that $h_{\alpha,\beta}(z_{\alpha,\beta}) = h(z)$; in other words, $h_{\alpha,\beta}(t) = h(\beta t + \alpha)$.

5.5. Normalized reductions. In many situations, we will also invoke the operation of taking a *normalized reduction* of a polynomial over $\bar{K}$, defined as follows.

Definition 1.8. A *normalized reduction* of a nonzero polynomial $h(z) \in \bar{K}[z]$ is the reduction in $k[z]$ of $\gamma^{-1}h$, where $\gamma \in \bar{K}^{\times}$ is some scalar satisfying $v(\gamma) = v(h)$.

Remark 1.9. Clearly a normalized reduction of a polynomial $h(z)$ is a nonzero polynomial in $k[x]$ and is unique up to scaling; thus, the degrees of the terms appearing in the normalized reduction (which is what we will be chiefly interested in for our purposes) do not depend on the particular choice of $\gamma \in K^{\times}$ in Definition 1.8.

CHAPTER 2

Semistable models of curves and their special fibers

The purpose of this section is to recall and develop definitions and results on semistable models of general curves over discretely-valued fields, which will later be applied to hyperelliptic curves.

1. Preliminaries on semistable models

In this subsection, we briefly recall a number of background results we will need about models of curves, for which our main reference will be [11]. In this section, C is a smooth, geometrically connected, projective curve over a complete discretely-valued field K, whose ring of integers is denoted R, and whose residue field k is assumed to be algebraically closed (see §1.5).

1.1. Curves and models. A *model* of C (over R) is a normal, flat, projective R-scheme $\mathcal{C}$ whose generic fiber is identified with C. The models of C form a preordered set Models(C), the order relation being given by *dominance*: given two models $\mathcal{C}$ and $\mathcal{C}'$ of C, we will write $\mathcal{C} \leq \mathcal{C}'$ to mean that $\mathcal{C}'$ dominates $\mathcal{C}$, i.e. that the identity $\mathrm{id} : C \to C$ extends to a birational morphism $\mathcal{C}' \to \mathcal{C}$. The preordered set Models(C) is filtered, meaning that given two models $\mathcal{C}_1$ and $\mathcal{C}_2$, it is always possible to find a model $\mathcal{C}$ dominating them both.

1.2. Special fibers of models and birational morphisms. The special fiber $\mathcal{C}_s$ of a model $\mathcal{C}$ of C is (geometrically) connected and consists of a number of irreducible components $V_1, \ldots, V_n$; these components are projective, possibly singular curves over the residue field k, each one appearing in $\mathcal{C}_s$ with a certain multiplicity (which is defined as the length of the local ring of $\mathcal{C}_s$ at the generic point of the component). We will denote by $\mathrm{Irr}(\mathcal{C}_s) = \{V_1, \ldots, V_n\}$ the set of such components. Since $\mathcal{C} \to \mathrm{Spec}(R)$ is proper and flat, the Euler-Poincaré characteristic of the generic fiber C and that of the special fiber $\mathcal{C}_s$ coincide; in other words, the genus $g(C)$ of the smooth K-curve C coincides with the arithmetic genus $p_a(\mathcal{C}_s)$ of the k-curve $\mathcal{C}_s$ (see, for example, [11, Proposition 8.3.28]).

When $\mathcal{C}$ and $\mathcal{C}'$ are two models such that $\mathcal{C} \leq \mathcal{C}'$, the image of a component V' of $\mathcal{C}'_s$ through the birational morphism $\mathcal{C}' \to \mathcal{C}$ is either a component V of $\mathcal{C}_s$ or a single point P of $\mathcal{C}_s$; in this second case, we say that the birational morphism $\mathcal{C}' \to \mathcal{C}$ *contracts* V'. The rule $V' \mapsto V$ defines a one-to-one correspondence between the irreducible components of $\mathcal{C}'_s$ that are not contracted by $\mathcal{C}' \to \mathcal{C}$ and the irreducible components of $\mathcal{C}_s$; we say that V is the *image* of V' in $\mathcal{C}_s$, and V' the *strict transform* of V in $\mathcal{C}'_s$. In other words, taking strict transforms defines an injection $\mathrm{Irr}(\mathcal{C}_s) \hookrightarrow \mathrm{Irr}(\mathcal{C}'_s)$, and $\mathrm{Irr}(\mathcal{C}'_s) \smallsetminus \mathrm{Irr}(\mathcal{C}_s)$ is the set of the irreducible components of $(\mathcal{C}')_s$ that the birational morphism contracts.

The birational morphism $\mathcal{C}' \to \mathcal{C}$ is an isomorphism precisely over the open subscheme $\mathcal{C} \smallsetminus \{P_1, \ldots, P_n\}$, where the P_i's are the points of the special fiber of $\mathcal{C}$ to which some $V' \in \mathrm{Irr}(\mathcal{C}'_s) \smallsetminus \mathrm{Irr}(\mathcal{C}_s)$ contracts. The fiber of $\mathcal{C}'_s$ above each P_i is connected of pure dimension 1, and its irreducible components are those components of $\mathcal{C}'_s$ that contract to P_i. If V is a component of $\mathcal{C}_s$ and V' is its strict transform in $\mathcal{C}'_s$, then the birational morphism of models $\mathcal{C}' \to \mathcal{C}$ restricts to a birational morphism of k-curves $V' \to V$.

1.3. Comparing models. Suppose that $\mathcal{C}$ and $\mathcal{C}'$ are two models of C. We can compare them by looking at the components of their special fibers. To this aim, let us make the auxiliary choice of a model $\mathcal{C}''$ dominating them both, so that we can think of $\mathrm{Irr}(\mathcal{C}_s)$ and $\mathrm{Irr}(\mathcal{C}'_s)$ as two subsets of a common larger set, namely $\mathrm{Irr}(\mathcal{C}''_s)$. We will denote by $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}') \subseteq \mathcal{C}_s(k)$ the set of points $P \in \mathcal{C}_s$ such that there exists an irreducible component $V'' \in \mathrm{Irr}(\mathcal{C}''_s)$ that is the strict transform of some component of $V' \in \mathrm{Irr}(\mathcal{C}'_s)$ and contracts to P. It is clear that the formation of $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}')$ does not depend on the choice of $\mathcal{C}''$.

Remark 2.1. In the language of [11, Subsection 8.3.2], $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}')$ is just the set of the centers in $\mathcal{C}$ of the R-valuations of $K(C)$ that are of the first kind in $\mathcal{C}'$ but not in $\mathcal{C}$.

Roughly speaking, this is the way that one should think of $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}')$: any given component $V \in \mathrm{Irr}(\mathcal{C}'_s)$ is either also present in the special fiber of $\mathcal{C}$ (i.e., $V \in \mathrm{Irr}(\mathcal{C}_s)$) or it is not, in which case it is *contracted* to some point $P_V \in \mathcal{C}_s(k)$. The set $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}')$ is simply the set of all such P_V's, as V varies in $\mathrm{Irr}(\mathcal{C}'_s) \smallsetminus \mathrm{Irr}(\mathcal{C}_s)$. We clearly have that $\mathrm{Ctr}(\mathcal{C}, \mathcal{C}') = \varnothing$ (i.e., all the irreducible components of $\mathcal{C}'_s$ are also present in $\mathcal{C}_s$) if and only if $\mathcal{C}$ dominates $\mathcal{C}'$.

1.4. Contracting components. Given a model $\mathcal{C}'$ of C and any proper subset $\{V_1, \ldots, V_n\} \subsetneq \mathrm{Irr}(\mathcal{C}'_s)$, it is always possible to form a model $\mathcal{C}$ of C dominated by $\mathcal{C}'$ such that the birational morphism $\mathcal{C}' \to \mathcal{C}$ contracts precisely the components $V_1, \ldots, V_n \in \mathrm{Irr}(\mathcal{C}'_s)$; as a consequence, we have $\mathrm{Irr}(\mathcal{C}_s) = \mathrm{Irr}(\mathcal{C}'_s) \smallsetminus \{V_1, \ldots, V_n\}$.

Given a finite number of models $\mathcal{C}_1, \ldots, \mathcal{C}_n$, one can form a minimal model $\mathcal{C}$ dominating them all: it is enough to take any model $\mathcal{C}'$ dominating them all, and then contract each $V \in \mathrm{Irr}(\mathcal{C}'_s)$ that is not the strict transform of an irreducible component of $(\mathcal{C}_i)_s$ for some i. It is clear that $\mathrm{Irr}(\mathcal{C}_s)$ coincides with the (non necessarily disjoint) union $\bigcup_i \mathrm{Irr}(\mathcal{C}_i)$.

1.5. Regular models. If we consider a model $\mathcal{C}$ of C that is regular (i.e. regular as an R-scheme), then it is possible to define the *intersection number* of any two components of $\mathcal{C}_s$; the resulting intersection matrix is negative semi-definite (see [11, Chapter 9]).

Definition 2.2. A component V of the special fiber of a regular model is said to be a *(-1)-line* if it is a line (i.e., $V \cong \mathbb{P}^1_k$) and its self-intersection number is -1. Similarly, it is said to be a (-2)-line if it is a line whose self-intersection number equals -2.

If one contracts any set of (-1)-lines in a regular model, it remains regular.

Given any model $\mathcal{C}$ of C, it is possible to find a regular model $\mathcal{C}'$ dominating $\mathcal{C}$. Moreover, among all regular models $\mathcal{C}'$ dominating $\mathcal{C}$, there is a minimum one (with respect to dominance), which is named the *minimal desingularization* of $\mathcal{C}$; it can be characterized as the unique regular model $\mathcal{C}'$ dominating $\mathcal{C}$ such that $\mathcal{C}' \to \mathcal{C}$ does not contract any (-1)-line of $\mathcal{C}'_s$. If $\mathcal{C}'$ is the minimal desingularization of $\mathcal{C}$, then the birational morphism $\mathcal{C}' \to \mathcal{C}$ fails to be an isomorphism precisely above the points of $\mathcal{C}_s$ at which $\mathcal{C}$ is not regular.

If the genus of C is positive, then, among all regular models of $\mathcal{C}$, there is a minimum one (with respect to dominance). This model is named the *minimal regular model* and will be denoted by $\mathcal{C}^{\min}$; it can be characterized as the unique regular model of C whose special fiber does not contain (-1)-lines.

1.6. Semistable models. A model $\mathcal{C}$ of C is said to be *semistable* if its special fiber is reduced and its singularities (if there are any) are all nodes (i.e. ordinary double points). More generally, we say that a model $\mathcal{C}$ of C is *semistable* at a point $P \in \mathcal{C}_s$ if $\mathcal{C}_s$ is reduced at P and if P is either a smooth point or a node of $\mathcal{C}_s$. Given a point $P \in \mathcal{C}_s$, if the model $\mathcal{C}$ is semistable at P then its completed local ring at P has the form $R[[t]]$, if $P \in \mathcal{C}_s$ is a smooth point, or $R[[t_1, t_2]]/(t_1t_2 - a)$ for some $a \in R$, with $v(a) > 0$ if $P \in \mathcal{C}_s$ is a node. The integer $v(a)/v(\pi) \geq 1$, where π is a uniformizer of R, is known as the *thickness* of the node. A semistable model is regular precisely when all of its nodes have thickness equal to 1.

To describe the combinatorics of a semistable model $\mathcal{C}$ of a curve C, one can form the *dual graph* $\Gamma(\mathcal{C}_s)$ of its special fiber, whose set of vertices is $\mathrm{Irr}(\mathcal{C}_s)$ and whose edges correspond to the nodes connecting them.

The notions of (-1)-line and (-2)-line given in Definition 2.2 for regular models can be extended to semistable ones as follows.

Definition 2.3. If $\mathcal{C}$ is a model, $V \in \mathrm{Irr}(\mathcal{C}_s)$, and $\mathcal{C}$ is semistable at the points of V, then V is said to be a (-1)-line (resp. a (-2)-line) if it is a line (i.e., $V \cong \mathbb{P}^1_k$) and the number of nodes of $\mathcal{C}_s$ lying on it is equal to 1 (resp. 2).

Remark 2.4. It is possible to show that, if $\mathcal{C}$ is regular model that is semistable at the points of a component $V \in \mathrm{Irr}(\mathcal{C}_s)$, then the definition above is consistent with the one given in Definition 2.2: this follows, for example, from the formula for self-intersection numbers given in [11, Proposition 9.1.21(b)].

Lemma 2.5. Suppose that $\mathcal{C}$ is a model of C that is semistable at the points of two components $V_1, V_2 \in \mathrm{Irr}(\mathcal{C}_s)$. If V_1 and V_2 are (-1)-lines, then they cannot intersect each other unless $g(C) = 0$.

Proof. Since V_1 and V_2 are (-1)-lines, if they intersect each other at a node, they cannot intersect any other irreducible components of $\mathcal{C}_s$. Since $\mathcal{C}_s$ is connected, this implies that $\mathcal{C}_s$ only consists of the two lines V_1 and V_2 crossing each other at a node, implying that $p_a(\mathcal{C}_s) = 0$ and consequently that $g(C) = 0$. □

Contracting (-1) and (-2)-lines does not ever disrupt semistability: more precisely, if $\mathcal{C}'$ is a model that is semistable at the points of some components $V_1, \ldots, V_n$ of $\mathcal{C}_s$, and the V_i's happen to all be (-1) and (-2)-lines, then the model $\mathcal{C}$ that is obtained from $\mathcal{C}$ by contracting all the V_i's is semistable at the points where the V_i's contract. Desingularizing is also an operation that preserves semistability: if $\mathcal{C}$ is semistable at a point $P \in \mathcal{C}_s$, and $\mathcal{C}'$ is its minimal desingularization, then $\mathcal{C}'$ is still semistable at all points lying above P; moreover, the desingularization $\mathcal{C}'$ is easy to describe:

(a) if P is a smooth point of $\mathcal{C}_s$, we have that $\mathcal{C}$ is regular at P, and the birational map $\mathcal{C}' \to \mathcal{C}$ is consequently an isomorphism above P; and
(b) if P is a node of thickness t, the inverse image of $\mathcal{C}'_s$ at P consist of a chain of t nodes of thickness 1, joined by $t - 1$ (-2)-lines (see [11, Section 5.3]).

More generally, if $\mathcal{C}$ is a model that is semistable at a point P and if $\mathcal{C}'$ is any model dominating $\mathcal{C}$ but dominated by its minimal desingularization, then $\mathcal{C}'$ is semistable at the points above P, and we have the following:

(a) if P is a smooth point of $\mathcal{C}_s$, then the birational map $\mathcal{C}' \to \mathcal{C}$ is an isomorphism above P;

(b) if P is a node of thickness t, then the inverse image of $\mathcal{C}'_s$ at P consist with a chain of $m \geq 1$ nodes whose thicknesses add up to t, joined by $m-1$ (-2)-lines (see [11, Section 5.3]).

When a semistable model exists, we say that C has *semistable reduction* over R. By Theorem 1.1 above, any curve C is guaranteed to have semistable reduction after replacing R with a large enough finite extension.

If C has semistable reduction and positive genus, its minimal regular model is semistable ([11, Theorem 10.3.34]). If C has semistable reduction and genus at least 2, then the set of its semistable models has a minimum (with respect to dominance), which is named the *stable model* of C; it is denoted by $\mathcal{C}^{\mathrm{st}}$ and can be characterized as the unique semistable model of C whose special fiber contains neither (-1)-lines nor (-2)-lines. The stable model $\mathcal{C}^{\mathrm{st}}$ can be obtained from $\mathcal{C}^{\mathrm{min}}$ by contracting all the (-2)-lines appearing in its special fiber.

1.7. Abelian, toric, and unipotent ranks. Given a model $\mathcal{C}$ of C, the Jacobian $\mathrm{Pic}^0(\mathcal{C}_s)$ of the (possibly singular) k-curve $\mathcal{C}_s$ is an extension of an abelian varitety A by a linear algebraic group, which in turn is an extension of a torus T by a smooth unipotent algebraic group U. The ranks of A, T, and U are respectively known as the abelian, toric, and unipotent rank of the special fiber $\mathcal{C}_s$ and will be denoted by $a(\mathcal{C}_s)$, $t(\mathcal{C}_s)$, and $u(\mathcal{C}_s)$ respectively: they are three non-negative integers adding up to the genus $p_a(\mathcal{C}_s) = g(C)$: see, for example, [11, Section 7.5] or [2, Chapters 8 and 9]. We have the following.

(a) For all models $\mathcal{C}$, the abelian rank $a(\mathcal{C}_s)$ coincides with the sum $a(\mathcal{C}_s) = \sum_{V \in \mathrm{Irr}(\mathcal{C}_s)} a(V)$, where $a(V) := g(\widetilde{V})$ is the genus of the normalization $\widetilde{V}$ of V.

(b) If $\mathcal{C}$ is a semistable model, the toric rank can be computed as $t(\mathcal{C}_s) = \dim_k H^1(\Gamma(\mathcal{C}_s))$, where $\Gamma(\mathcal{C}_s)$ is the dual graph of $\mathcal{C}_s$ (see [2, Example 9.8] for a proof); in other words, we have $t(\mathcal{C}_s) = N_{\mathrm{nodes}}(\mathcal{C}_s) - N_{\mathrm{irr}}(\mathcal{C}_s) + 1$, where $N_{\mathrm{nodes}}(\mathcal{C}_s)$ denotes the number of nodes, and $N_{\mathrm{irr}}(\mathcal{C}_s)$ is the number of irreducible components (i.e., the cardinality of $\mathrm{Irr}(\mathcal{C}_s)$).

(c) The unipotent rank $u(\mathcal{C}_s)$ is 0 if $\mathcal{C}$ is a semistable model.

Under certain hypotheses, the abelian, unipotent and toric rank do not depend on the chosen model. In particular, we have the following.

Proposition 2.6. If $\mathcal{C}$ and $\mathcal{C}'$ are two models of C over R, and if each of these models is either regular or semistable, then we have $a(\mathcal{C}_s) = a(\mathcal{C}'_s)$, $t(\mathcal{C}_s) = t(\mathcal{C}'_s)$, and $u(\mathcal{C}_s) = u(\mathcal{C}'_s)$.

Proof. It is clearly enough to prove the result (a) when $\mathcal{C}$ and $\mathcal{C}'$ are both regular, and (b) when $\mathcal{C}$ is a semistable model and $\mathcal{C}'$ is its minimal desingularization. In case (a), one can form a regular model $\mathcal{C}''$ dominating them both and apply [11, Lemma 10.3.40]. In case (b), it follows from the description given in §2.1.1.6 of the desingularization of a semistable model that $\mathcal{C}'$ is also semistable, and we get $N_{\mathrm{nodes}}(\mathcal{C}_s) - N_{\mathrm{irr}}(\mathcal{C}_s) = N_{\mathrm{nodes}}(\mathcal{C}'_s) - N_{\mathrm{irr}}(\mathcal{C}'_s)$, which is to say $t(\mathcal{C}_s) = t(\mathcal{C}'_s)$ Moreover, since any new components introduced by the desingularization process are lines, we have $a(\mathcal{C}_s) = a(\mathcal{C}'_s)$, and finally, since both $\mathcal{C}$ and $\mathcal{C}'$ are semistable, we have $u(\mathcal{C}_s) = u(\mathcal{C}'_s) = 0$. $\square$

1.8. Field extensions. Suppose we are given a finite extension K'/K (which, under our assumptions, will necessarily be totally ramified, since the residue field k of K is algebraically closed); let $e_{K'/K}$ be the ramification index (which coincides, in our setting, with the degree of the extension), and let $R' \supseteq R$ denote the ring of integers of K'. We will freely say a *model of C over R'* to mean a model of $C' := C \otimes_K K'$ over R' as defined in §2.1.1.1. Given a model $\mathcal{C}$ of C over R, it is possible to construct a corresponding model $\mathcal{C}'$ of C over R', which is defined as the normalization of the base-change $\mathcal{C} \otimes_R R'$. When $\mathcal{C}$ has reduced special fiber (and hence, in particular, when $\mathcal{C}$ is semistable), the scheme $\mathcal{C} \otimes_R R'$ is already normal (for example, by Serre's criterion for normality), and so we have $\mathcal{C}' = \mathcal{C} \otimes_R R'$: in this last case, the special fibers $\mathcal{C}'_s$ and $\mathcal{C}_s$ are canonically isomorphic.

We remark that regularity is not preserved in general when R gets extended: if $\mathcal{C}$ is a regular model over R, the corresponding model $\mathcal{C}'$ over an extension R' may no longer be regular. Semistability, however, is preserved: whenever $\mathcal{C}$ is semistable, $\mathcal{C}'$ is semistable too; however, the thickness of each node of $\mathcal{C}$ gets multiplied by the ramification index $e_{K'/K}$ in $\mathcal{C}'$.

In this section we look more closely at the special fibers of models (over R) of a smooth projective geometrically connected K-curve C. In §2.2, in particular, we define a number of invariants attached to each component of the special fiber of a model, while in §2.3, we use them to state a criterion that allows us to identify those models of C that are part of the minimal regular model when C has semistable reduction.

2. Invariants attached to a component of the special fiber

Given a model $\mathcal{C}$ of C and a component $V \in \mathrm{Irr}(\mathcal{C}_s)$, we consider several invariants attached to V, listed as follows:

(1) $m(V)$ denotes the multiplicity of V in $\mathcal{C}_s$;
(2) $a(V)$ denotes the *abelian rank* of V, i.e. the genus of the normalization $\widetilde{V}$ of V;
(3) $w(V)$ is defined only when $m(V) = 1$, and it denotes the number of singular points of $\mathcal{C}_s$ that belong to V, each one counted as many times as the number of branches of V at that point; in other words, if $\tilde{V}$ is the normalization of V, then $w(V)$ is the number of points of $\tilde{V}$ that lie over $V \cap \mathrm{Sing}(\mathcal{C}_s)$, where $\mathrm{Sing}(\mathcal{C}_s)$ is the set of singular points of $\mathcal{C}_s$.

We now show that, under appropriate assumptions, the integers m, a, and w are left invariant when the model is changed.

Lemma 2.7. Let $\mathcal{C}'$ be another model of C which dominates $\mathcal{C}$, and let V' denote the strict transform of V in $\mathcal{C}'_s$. Then we have $m(V') = m(V)$ and $a(V') = a(V)$. Moreover, if $\mathcal{C}'$ is dominated by the minimal desingularization of $\mathcal{C}$, we also have $w(V') = w(V)$.

Proof. For m and a, the lemma immediately follows from the consideration that $\mathcal{C}'_s \to \mathcal{C}_s$ is an isomorphism away from a finite set of points of $\mathcal{C}_s$. We will now prove the result for w.

Let Q be a point of V' which lies over some $P \in V$. We claim that $\mathcal{C}'_s$ is smooth (resp. singular) at Q if and only if $\mathcal{C}_s$ is smooth (resp. singular) at P. This is obvious whenever $\mathcal{C}' \to \mathcal{C}$ is an isomorphism above P. If $\mathcal{C}' \to \mathcal{C}$ is not an isomorphism above P, the claim follows from the two following observations. Firstly, since we are assuming that $\mathcal{C}'$ is dominated by the minimal desingularization of $\mathcal{C}$, it must be the case that $\mathcal{C}$ is not regular at P, and consequently that $\mathcal{C}_s$ is singular at P. Secondly, the fiber of $\mathcal{C}' \to \mathcal{C}$ above P is pure of dimension 1, and it consists of those components E_i of $\mathcal{C}'_s$ that contract to P; the point Q will thus belong not only to V', but also to one of the E_i's, so that $\mathcal{C}'_s$ will certainly be singular at Q. This completes the proof of the claim.

Now, since $\mathcal{C}' \to \mathcal{C}$ restricts to a birational morphism $V' \to V$ of k-curves, the set of branches of V at a point $P \in \mathcal{C}_s$ equals the set of branches of V' at the points of $\mathcal{C}'_s$ lying above P. If we combine this consideration with the claim we have just proved, we have that $V' \to V$ induces a bijection between the set $\mathfrak{B}$ of the branches of V at the singular points of $\mathcal{C}_s$ and the set $\mathfrak{B}'$ of the branches of V' at the singular points of $\mathcal{C}'_s$. The equality $w(V') = w(V)$ follows. □

We now describe how the invariants we have defined allow us to detect (-1)-lines and (-2)-lines.

Lemma 2.8. Let $\mathcal{C}$ be any model of C, and let V be an irreducible component of $\mathcal{C}_s$. Then,

(a) if $\mathcal{C}$ is regular and V is a (-1)-line of multiplicity 1, then $a(V) = 0$ and $w(V) = 1$;
(b) if $\mathcal{C}$ is regular and V is a (-2)-line of multiplicity 1, then $a(V) = 0$ and $w(V) \in \{1, 2\}$;
(c) if $\mathcal{C}$ is semistable at the points of V, then V is a (-1)-line if and only if $a(V) = 0$ and $w(V) = 1$; and
(d) if $\mathcal{C}$ is semistable at the points of V, then V is a (-2)-line if and only if $a(V) = 0$ and $w(V) = 2$ (the reverse implication only holds if $g(C) \geq 2$).

Proof. If V is a component of multiplicity 1 in the special fiber $\mathcal{C}_s$ of a regular model $\mathcal{C}$, then it follows from the intersection theory of regular models (see [11, Chapter 9]) that its self-intersection number of V is equal to minus the number of points at which V intersects the other components of $\mathcal{C}_s$, each counted with a certain (positive) multiplicity. Once this has been observed, parts (a) and (b) follow immediately from the definition of (-1)-lines and (-2)-lines for regular models (Definition 2.2).

Suppose now that V is a component of the special fiber $\mathcal{C}_s$ of a model $\mathcal{C}$ that is semistable at the points of V (which, in particular, implies $m(V) = 1$). From the definition of the invariant w and the structure of semistable models, it is clear that $w(V)$ equals the sum $2w_{\mathrm{self}}(V) + w_{\mathrm{other}}(V)$, where $w_{\mathrm{self}}(V)$ is the number of self-intersections of V, while $w_{\mathrm{other}}(V)$ is the number of intersections of V with other components of $\mathcal{C}_s$; moreover, we have $w_{\mathrm{self}}(V) = 0$ if and only if V is smooth. But the line $\mathbb{P}^1_k$ is the unique smooth k-curve with abelian rank 0, so the component V is a line if and only if $a(V) = 0$ and $w_{\mathrm{self}}(V) = 0$; according to Definition 2.3, the component V is thus a (-1)-line or a (-2)-curve if and only if $a(V) = 0$, $w_{\mathrm{self}}(V) = 0$, and $w_{\mathrm{other}}(V)$ equals 1 or 2 respectively.

From the considerations above, both implications of (c), as well the forward implication of (d), immediately follow. To prove the reverse implication of (d), one has to exclude the possibility that $a(V) = 0$, $w_{\mathrm{self}}(V) = 1$, and $w_{\mathrm{other}}(V) = 0$. But if this were the case, the unique irreducible component of $\mathcal{C}_s$ would be V (because $w_{\text{other}}(V) = 0$, but $\mathcal{C}_s$ is connected), and the special fiber $\mathcal{C}_s$ would consequently be a reduced k-curve having arithmetic genus equal to that of V, which is $a(V) + w_{\text{self}}(V) = 1$. Since the arithmetic genus of $\mathcal{C}_s$ coincides with $g(C)$, we would get $g(C) = 1$; we therefore get the reverse implication of (d) as long as $g(C) \neq 1$. □

Inspired by the above lemma, we make the following definition.

Definition 2.9. Given a model $\mathcal{C}$ and an irreducible component V of $\mathcal{C}_s$ at whose points $\mathcal{C}$ is semistable, the component V is said to be a *(-2)-curve* of $\mathcal{C}_s$ if $m(V) = 1$, $a(V) = 0$ and $w(V) = 2$.

Remark 2.10. Lemma 2.8 ensures that, if V is a component of $\mathcal{C}_s$ and $\mathcal{C}$ is semistable at the points of V, then, when V is (-2)-line, it is a (-2)-curve, and the converse also holds provided that $g(C) \neq 1$. If $g(C) = 1$, the proof of Lemma 2.8 shows that V may be a (-2)-curve without being a (-2)-line, and this happens precisely when V is the unique component of $\mathcal{C}_s$ and it is a k-curve of abelian rank 0 intersecting itself once (which is to say, a projective line with two points identified).

The properties of being a (-1)-line or a (-2)-curve are preserved and reflected under desingularization in the semistable case; this is the reason why the notion of a (-2)-curve (rather than a (-2)-line) will turn out to be more convenient for us.

Proposition 2.11. Let $\mathcal{C}$ be a model; let V be a component of $\mathcal{C}_s$; and let $\mathcal{C}'$ be a model dominating $\mathcal{C}$ but dominated by the minimal desingularization of $\mathcal{C}$. Assume that $\mathcal{C}$ is semistable at the points of V. Let V' denote the strict transform of V in $\mathcal{C}'_s$. We have that $\mathcal{C}'$ is semistable at the points of V', and V' is a (-1)-line (resp. a (-2)-curve) if and only if V is.

Proof. The fact that $\mathcal{C}'$ is semistable at the points of V' has been discussed in §2.1.1.6. We have seen how, in the present setting, being a (-1)-line or a (-2)-curve is something that can be characterized by means of the invariants a and w. Hence, the result follows from Lemma 2.7. □

3. A criterion for being part of the minimal regular model

As initial evidence of the usefulness of the invariants a, m, w introduced before, we provide a criterion for a model $\mathcal{C}$ to be part of the minimal regular model $\mathcal{C}^{\min}$ in the case that C has semistable reduction.

Proposition 2.12. Assume that $g(C) \geq 1$ and that C has semistable reduction. Let $\mathcal{C}$ be any model. Then $\mathcal{C} \leq \mathcal{C}^{\min}$ if and only if for each component V of $\mathcal{C}_s$, we have

(i) $m(V) = 1$, and
(ii) either $a(V) \geq 1$ or $w(V) \geq 2$.

Proof. First assume that we have $\mathcal{C} \leq \mathcal{C}^{\min}$. Then the invariants a, m, and w of a vertical component V of $\mathcal{C}$ must be equal to those of its strict transform $V^{\min}$ in $\mathcal{C}^{\min}$, thanks to Lemma 2.7 (more generally, they remain the same in any model lying between $\mathcal{C} \leq \mathcal{C}^{\min}$). Since $\mathcal{C}^{\min}$ is semistable, its special fiber is reduced; thus, we get $m(V) = 1$. Let us now assume that $a(V) = 0$. If it were the case that $w(V) = 0$, then $\mathcal{C}_s = V$

would be a line; since the arithmetic genus of $\mathcal{C}_s$ coincides with $g(C)$, this contradicts the condition that $g(C) \geq 1$. If we had $w(V) = 1$, then, via Lemma 2.8(c), $V^{\min}$ would be a (-1)-line, which is impossible, since the minimal regular model does not contain (-1)-lines. Thus, the quantity $w(V)$ is necessarily at least 2.

Now assume that for each component V of $\mathcal{C}_s$, the conditions (i) and (ii) given in the statement hold. Let $\mathcal{C}'$ be the minimal desingularization of $\mathcal{C}$. Assume by way contradiction that $\mathcal{C}'_s$ contains a (-1)-line. Since the desingularization $\mathcal{C}'$ is minimal, such a (-1)-line must necessarily be the strict transform V' of some component $V \in \mathrm{Irr}(\mathcal{C}_s)$. By Lemma 2.7, the quantities $a(V')$, $m(V')$ and $w(V')$ are equal to $a(V)$, $m(V)$ and $w(V)$ respectively. Thus, from the condition $m(V) = 1$, we deduce $m(V') = 1$; since V' is a (-1)-line of multiplicity 1, Lemma 2.8(a) ensures that $a(V') = 0$ and $w(V') = 1$, hence $a(V) = 0$ and $w(V) = 1$. But this contradicts our hypothesis, so we conclude that $\mathcal{C}'_s$ cannot contain a (-1)-line. It follows that $\mathcal{C}' = \mathcal{C}^{\min}$, and we get $\mathcal{C} \leq \mathcal{C}^{\min}$ as desired. $\square$

CHAPTER 3

The relatively stable model

In this section, we assume that $Y \to X$ is a Galois cover of smooth projective geometrically connected curves over K; let $G := \mathrm{Aut}_X(Y) = \mathrm{Aut}_{K(X)}(K(Y))$ denote the Galois group. We remark that, by the Riemann-Hurwitz formula, we always have that $g(Y) \geq g(X)$.

In §3.1 we collect some well-known background results about semistable models of Galois covers, for which good references are [11] and [12]. In §3.2 and in §3.3, we study nodes and (-2)-curves of a semistable model of Y with respect to the cover $Y \to X$: these preliminaries allow us to define, in §3.4, a particular semistable model of Y (relative to the Galois cover $Y \to X$) that we name the *relatively stable model* of Y; we denote it $\mathcal{Y}^{\mathrm{rst}}$, and it arises as the normalization in $K(Y)$ of a semistable model $\mathcal{X}^{(\mathrm{rst})}$ of the line X. Existence and uniqueness results for $\mathcal{Y}^{\mathrm{rst}}$ hold provided that we have $g(Y) \geq 2$ or that we have $g(Y) = 1$ and $g(X) = 0$ (existence is only guaranteed if one allows replacing R with a large enough extension). In §3.5 we will explain some methods for detecting the components of $\mathcal{Y}^{\mathrm{rst}}$.

1. Models of Galois covers

In the setting describe above, one can produce models of Y from models of X. More precisely, to each model $\mathcal{X}$ of X we can attach a corresponding model $\mathcal{Y}$ of Y by taking the normalization of $\mathcal{X}$ in the function field $K(Y)$ – we will say that $\mathcal{Y}$ *comes from* $\mathcal{X}$ (or that $\mathcal{Y}$ is the model of Y *corresponding to* $\mathcal{X}$). Given two models $\mathcal{X}$ and $\mathcal{X}'$ of X, if $\mathcal{Y}$ and $\mathcal{Y}'$ are the corresponding models of Y, then it is easy to show that $\mathcal{X} \leq \mathcal{X}'$ if and only if $\mathcal{Y} \leq \mathcal{Y}'$: in other words, normalizing in $K(Y)$ defines an embedding of preordered sets $\mathrm{Models}(X) \hookrightarrow \mathrm{Models}(Y)$. The essential image of the embedding consists of those models $\mathcal{Y}$ of Y on which G acts, i.e. those for which the action of G on the generic fiber Y extends (in a necessarily unique way) to an action on the R-scheme $\mathcal{Y}$. Given a model $\mathcal{Y}$ of Y on which G acts, the model of X from which $\mathcal{Y}$ comes can be recovered as the quotient $\mathcal{Y}/G$.

If $\mathcal{X}$ is a model of X and $\mathcal{Y}$ is the corresponding model of Y, the set of irreducible components $\mathrm{Irr}(\mathcal{Y}_s)$ is a G-set, and we have $\mathrm{Irr}(\mathcal{X}_s) = \mathrm{Irr}(\mathcal{Y}_s)/G$. Given $\mathcal{X}, \mathcal{X}' \in \mathrm{Models}(X)$ and letting $\mathcal{Y}, \mathcal{Y}'$ be the corresponding models of Y, it is also not difficult to see that $\mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}') = f^{-1}(\mathrm{Ctr}(\mathcal{X}, \mathcal{X}'))$ (see §2.1.1.3 for notation), where $f : \mathcal{Y} \to \mathcal{X}$ is the cover map.

The minimal regular model $\mathcal{Y}^{\mathrm{min}}$ and the stable model $\mathcal{Y}^{\mathrm{st}}$, when defined, are always acted upon by G; we will use the notation $\mathcal{X}^{(\mathrm{min})} = \mathcal{Y}^{\mathrm{min}}/G$ and $\mathcal{X}^{\mathrm{st}} = \mathcal{Y}^{\mathrm{st}}/G$ to denote the models of X from which they come.

2. Vertical and horizontal (-2)-curves

Given a model $\mathcal{Y}$ of Y coming from a model $\mathcal{X}$ of X and a component $V \in \operatorname{Irr}(\mathcal{Y}_s)$ of multiplicity 1, we replace the invariant $w := w(V)$ introduced in §2.2 with a richer datum $\underline{w} := \underline{w}(V)$ that takes into account the action of G. We have that G acts on $\operatorname{Irr}(\mathcal{Y}_s)$, and we denote by G_V the stabilizer of V with respect to this action. If $\mathfrak{B} := \{b_1, \dots, b_w\}$ are the branches of V passing through the singular points of $\mathcal{Y}_s$, the stabilizer G_V clearly acts on $\mathfrak{B}$, and we define $\underline{w}(V)$ to be the partition of the integer $w(V) = |\mathfrak{B}|$ given by the cardinality of the orbits of the G_V-set $\mathfrak{B}$.

Lemma 3.1. If $\mathcal{Y}'$ is a model acted upon by G which dominates $\mathcal{Y}$ and is dominated by the minimal desingularization of $\mathcal{Y}$, then we have $\underline{w}(V') = \underline{w}(V)$, where V' is the strict transform of V in $(\mathcal{Y}')_s$.

Proof. The set $\mathfrak{B}$ of the branches of V passing through the singular points of $\mathcal{Y}_s$ does not change as we replace $\mathcal{Y}$ with $\mathcal{Y}'$, and V with its strict transform, as was shown in the proof of Lemma 2.7; since the birational map $\mathcal{Y}' \to \mathcal{Y}$ is G-equivariant, the stabilizers G_V and $G_{V'}$ coincide, and $\mathfrak{B}$ is preserved not only as a set, but also as a G_V-set. □

We recall that, when $\mathcal{Y}$ is semistable at the points of some component $V \in \operatorname{Irr}(\mathcal{Y}_s)$, we say that V is a (-2)-curve whenever $a(V) = 0$ and $w(V) = 2$ (see Definition 2.9). In our setting, the presence of a G-action on $\mathcal{Y}$ allow us to distinguish between *vertical* and *horizontal* (-2)-curves, according to the two possible values for the invariant $\underline{w}(V)$.

Definition 3.2. Given a component $V \in \operatorname{Irr}(\mathcal{Y}_s)$ such that $\mathcal{Y}$ is semistable at the points of V, if V is a (-2)-curve we say that it is *vertical* or *horizontal* depending on whether $\underline{w}(V) = (2)$ or $\underline{w}(V) = (1, 1)$.

The property of being a horizontal or vertical (-2)-line is preserved and reflected under desingularization of semistable models.

Proposition 3.3. Let $\mathcal{Y}'$ another model acted upon by G which dominates $\mathcal{Y}$ but is dominated by the minimal desingularization of $\mathcal{Y}$. Given V a component of $\mathcal{Y}_s$ such that $\mathcal{Y}$ is semistable at the points of V, if V' denotes the strict transform of V in $\mathcal{Y}'_s$, we have that $\mathcal{Y}'$ is semistable at the points of V'; moreover, the transform V' is a (-1)-line (resp. a horizontal (-2)-curve, resp. a vertical (-2)-curve) if and only if V is.

Proof. The proof is analogous to that of Proposition 2.11, taking into account Lemma 3.1. □

3. Vanishing and persistent nodes

Given a model $\mathcal{Y}$ of Y corresponding to some model $\mathcal{X}$ of X, we can ask ourselves how the properties of $\mathcal{X}$ and $\mathcal{Y}$ are related to each other. If we write $f : \mathcal{Y} \to \mathcal{X} = \mathcal{Y}/G$ for the quotient map, we present an important result concerning the semistability of $\mathcal{X}$ and $\mathcal{Y}$.

Proposition 3.4. In the setting above, we have that $\mathcal{X}$ is semistable at $f(Q)$ whenever $\mathcal{Y}$ is semistable at some $Q \in \mathcal{Y}_s$. More precisely, we have the following.

(a) If Q is a smooth point of $\mathcal{Y}_s$, then $f(Q)$ is a smooth point of $\mathcal{X}_s$;
(b) If Q is a node of thickness t of $\mathcal{Y}_s$, we have two possibilities:

(i) if the stabilizer $G_Q \leq G$ of Q permutes the two branches of $\mathcal{Y}_s$ passing through Q, then $f(Q)$ is a smooth point of $\mathcal{X}_s$;
(ii) if, instead, the stabilizer $G_Q \leq G$ of Q does not flip the two branches of $\mathcal{Y}_s$ passing through Q, then $f(Q)$ is a node of $\mathcal{X}_s$, and its thickness is $t|G_Q|$.

In particular, if the model $\mathcal{Y}$ is semistable, then so is the model $\mathcal{X}$.

Proof. The proof consists of an explicit local study of the quotient map $f : \mathcal{Y} \to \mathcal{X}$, which can be found in [11, Proposition 10.3.48]. $\square$

Definition 3.5. A node Q of $\mathcal{Y}$ is said to be *vanishing* or *persistent* with respect to the Galois cover $Y \to X$ depending on whether it falls under case (i) or (ii) of Proposition 3.4(b), i.e. depending on whether it lies above a smooth point or a node of $\mathcal{X}_s$.

Remark 3.6. We have already remarked in §3.1 that $f : \mathcal{Y} \to \mathcal{X}$ induces a one-to-one correspondence between the irreducible components of $\mathcal{X}_s$ and the G-orbits of irreducible components of $\mathcal{Y}_s$. Proposition 3.4 and Definition 3.5 tell us that the nodes of $\mathcal{X}$ correspond to the G-orbits of *persistent* nodes of $\mathcal{Y}$.

We have seen in §2.1.1.6 that, if Q is a node of thickness t of $\mathcal{Y}$, its inverse image in the special fiber of the minimal desingularization of $\mathcal{Y}$ consists of a chain of t nodes of thickness 1, connected by $t-1$ (-2)-curves. More generally, if $\mathcal{Y}'$ is any model acted upon by G that dominates $\mathcal{Y}$ but is dominated by its minimal desingularization, then $\mathcal{Y}'$ is semistable at the points lying above Q, and the inverse image of Q in $\mathcal{Y}'_s$ consists of a chain of m nodes $Q_1, \ldots, Q_m$, whose thicknesses add up to t, and $m-1$ (-2)-curves $L_1, \ldots, L_{m-1}$ connecting them. It is an interesting question to ask whether the Q_i's are persistent or vanishing, and whether the L_i's are horizontal or vertical.

Proposition 3.7. In the setting above, we have the following:

(a) if $Q \in \mathcal{Y}_s$ is a persistent node, then the Q_i's also are, and the L_i's are all horizontal;
(b) if $Q \in \mathcal{Y}_s$ is vanishing and m is odd, then G_Q permutes L_i and L_{m-i}; the L_i's are all horizontal, while the Q_i's are all persistent, apart from the middle one $Q_{(m+1)/2}$ which is vanishing; and
(c) if $Q \in \mathcal{Y}_s$ is vanishing and m is even, then G_Q permutes each L_i with L_{m-i}; the L_i's are all horizontal, apart from the middle one $L_{m/2}$ which is vertical, while the Q_i's are all persistent.

Proof. We remark that, since the Q_i's and the L_i's have image Q in $\mathcal{Y}_s$, we have $G_{Q_i} \leq G_Q$ and $G_{L_i} \leq G_Q$, and every $g \in G_Q$ acts on the set of the sets $\{L_i\}_{1 \leq i \leq m-1}$ and $\{Q_i\}_{1 \leq i \leq m}$. We denote by Λ_- and Λ_+ the strict transforms in $\mathcal{Y}'_s$ of the two branches of $\mathcal{Y}_s$ passing through Q, so that the node Q_1 connects Λ_- with L_1 and the node Q_m connects L_{m-1} with Λ_+.

Choose an element $g \in G_Q$ which fixes Λ_+ and Λ_-. Since Q_1 is the unique point of Λ_- lying above Q, the point Q_1 is also fixed by g; since g fixes Q_1 and Λ_-, it must also fix the only other branch of $\mathcal{Y}_s$ passing through Q_1; therefore, it fixes L_1. Since g fixes Q_1 and L_1, it must also fix the only other node that lies on L_1, namely Q_2. Iterating the argument, one gets that the Galois element g stabilizes each of the L_i's and the Q_i's.

Now choose an element $g \in G_Q$ that flips Λ_+ and Λ_-. Since Q_1 is a point of Λ_-, the point $g \cdot Q_1$ must belong to Λ_+ and lie above Q; we therefore have $g \cdot Q_1 = Q_m$. From the fact that $g \cdot Q_1 = Q_m$ and $g \cdot \Lambda_- = \Lambda_+$, one deduces that $g \cdot L_1 = L_{m-1}$, and so on. From this kind of iterative argument, the results follow (how it ends clearly depends on whether m is even or odd). $\square$

4. Defining the relatively stable model

We are now ready to define the *relatively stable model* of Y with respect to the Galois cover $Y \to X$.

Definition 3.8. A model of $\mathcal{Y}$ of Y is said to be *relatively stable* with respect to the Galois cover $Y \to X$ if it is semistable, it is acted upon by G, and its special fiber does not contain vanishing nodes, (-1)-lines, or horizontal (-2)-curves. If a relatively stable model exists, the curve Y is said to have *relatively stable reduction* with respect to the cover $Y \to X$.

Remark 3.9. If $\mathcal{Y}$ is relatively stable and $\mathcal{X} = \mathcal{Y}/G$, since $\mathcal{Y}$ cannot contain vanishing nodes, we have $\mathrm{Sing}(\mathcal{Y}_s) = f^{-1}(\mathrm{Sing}(\mathcal{X}_s))$, where $f : \mathcal{Y} \to \mathcal{X}$ is the cover map, while $\mathrm{Sing}(\mathcal{X}_s)$ and $\mathrm{Sing}(\mathcal{Y}_s)$ are the respective sets of nodes of the semistable models $\mathcal{X}$ and $\mathcal{Y}$.

A relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ can only exist if the curve Y has semistable reduction; moreover, since $\mathcal{Y}^{\mathrm{rst}}$ is semistable and contains no (-1)-lines, it is clear that $\mathcal{Y}^{\mathrm{rst}} \leq \mathcal{Y}^{\mathrm{min}}$ (provided that $\mathcal{Y}^{\mathrm{min}}$ exists, i.e. $g(Y) \geq 1$). It is also clear from the definition that the property of being relatively stable is preserved and reflected under arbitrary extensions of R. Finally, we observe that, if the cover $Y \to X$ is trivial, a relatively stable model of Y is nothing but a stable model.

Proposition 3.10. Assume that $g(Y) \geq 1$. The relatively stable model, if it exists, is unique.

Proof. Let $\mathcal{Y}_1$ and $\mathcal{Y}_2$ two relatively stable models, and let $\mathcal{Y}_3$ be the minimal model dominating them both; by minimality, each vertical component of $(\mathcal{Y}_3)_s$ is the strict transform of a component of $(\mathcal{Y}_1)_s$ or of a component of $(\mathcal{Y}_2)_s$. Since $\mathcal{Y}_1$ and $\mathcal{Y}_2$ are $\leq \mathcal{Y}^{\mathrm{min}}$, we also have $\mathcal{Y}_3 \leq \mathcal{Y}^{\mathrm{min}}$; moreover, the model $\mathcal{Y}_3$ is semistable since it lies between the semistable model $\mathcal{Y}_1$ and its minimal desingularization $\mathcal{Y}^{\mathrm{min}}$. Since the special fibers of $\mathcal{Y}_1$ and $\mathcal{Y}_2$ only contain vertical (-2)-curves, the same is true for $\mathcal{Y}_3$, thanks to Proposition 3.3. Suppose by way of contradiction that $\mathcal{Y}_3 \gneq \mathcal{Y}_1$: this means that some node Q of $(\mathcal{Y}_1)_s$ is replaced, in $(\mathcal{Y}_3)_s$, by a chain of m nodes and $m-1$ (-2)-curves (with $m > 1$). But since $(\mathcal{Y}_3)_s$ does not contain horizontal (-2)-curves, Proposition 3.7 forces Q to be vanishing of thickness 2, which is a contradiction, since $(\mathcal{Y}_1)_s$ does not contain vanishing nodes. Hence, we have $\mathcal{Y}_1 = \mathcal{Y}_3$, which is to say that $\mathcal{Y}_1 \geq \mathcal{Y}_2$; now we get $\mathcal{Y}_1 = \mathcal{Y}_2$ by symmetry. $\square$

From now on, we use the symbol $\mathcal{Y}^{\mathrm{rst}}$ to denote the relatively stable model of Y (whenever it exists), while $\mathcal{X}^{(\mathrm{rst})} = \mathcal{Y}^{\mathrm{rst}}/G$ denotes the model of the line X to which it corresponds in the sense of §3.1.

Lemma 3.11. Assume that $g(Y) \geq 1$ and that Y has a semistable model $\mathcal{Y}$ acted upon by G whose special fiber only consists of horizontal (-2)-curves connected by persistent nodes. Then we have $g(Y) = g(X) = 1$.

PROOF. Let $\mathcal{X} = \mathcal{Y}/G$ be the semistable model of X from which $\mathcal{Y}$ comes. Since all irreducible components of $\mathcal{Y}_s$ have abelian rank 0, the same is also be true for all irreducible components of $\mathcal{X}_s = \mathcal{Y}_s/G$; if a denotes the abelian rank, we thus have that $a(\mathcal{Y}_s) = a(\mathcal{X}_s) = 0$.

Since $\mathcal{Y}_s$ only consists of (-2)-curves, we have that its dual graph $\Gamma(\mathcal{Y}_s)$ is a polygon with $N \geq 1$ sides (when $N = 1$, it consists of a single vertex, and a loop around it). We clearly have an action of G on $\Gamma(\mathcal{Y}_s)$; moreover, the absence of vanishing nodes ensures that $\Gamma(\mathcal{X}_s) = \Gamma(\mathcal{Y}_s)/G$ (see Remark 3.6).

Let σ_g be the automorphism of the polygon $\Gamma(\mathcal{Y}_s)$ induced by an element $g \in G$. Then, σ_g cannot be a reflection: this is because a reflection either fixes a vertex and flips the two edges it connects, or it fixes an edge flipping the two vertices lying on it; in the first case, $\mathcal{Y}_s$ would contain a vanishing node, and in the second one it would contain a vertical (-2)-curve. Hence, the automorphism σ_g is necessarily a rotation; the image of G in $\mathrm{Aut}(\mathcal{G})$ is consequently a cyclic subgroup consisting of d rotations for some $d|N$, and $\Gamma(\mathcal{X}_s) = \Gamma(\mathcal{Y}_s)/G$ is consequently a polygon with N/d sides. If t denotes the toric rank, we consequently have $t(\mathcal{X}_s) = t(\mathcal{Y}_s) = 1$.

By the results in §2.1.1.7, we can now conclude that $g(Y) = g(X) = 1$. □

Proposition 3.12. Assume that $g(Y) \geq 2$, or that $g(Y) = 1$ and $g(X) = 0$. The relatively stable model exists if and only if Y has semistable reduction and $\mathcal{Y}^{\min}$ contains no vanishing nodes. If it does exist, the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ can be formed by contracting all horizontal (-2)-curves of the minimal regular model $\mathcal{Y}^{\min}$.

PROOF. Suppose the curve Y has semistable reduction and that $\mathcal{Y}^{\min}$ contains no vanishing nodes. By Lemma 3.11, the special fiber $(\mathcal{Y}^{\min})_s$ cannot consist only of horizontal (-2)-curves, as that would contradict our hypothesis on genera.

We are consequently allowed to form a model $\mathcal{Y}^{\mathrm{rst}}$ by contracting all horizontal (-2)-curves of $\mathcal{Y}^{\min}$, and it will still be semistable. It is clear that G acts on $\mathcal{Y}^{\mathrm{rst}}$. Suppose that $(\mathcal{Y}^{\mathrm{rst}})_s$ contains a horizontal (-2)-curve. Then its strict transform $(\mathcal{Y}^{\min})_s$ is still a horizontal (-2)-curve in light of Proposition 3.3, which contradicts the fact that all (-2)-curves of $(\mathcal{Y}^{\min})_s$, by construction, get contracted in $(\mathcal{Y}^{\mathrm{rst}})_s$.

Suppose now that $(\mathcal{Y}^{\mathrm{rst}})_s$ contains a vanishing node. Then, in light of Proposition 3.7, its preimage in $(\mathcal{Y}^{\min})_s$ must contain a vanishing node or a vertical (-2)-curve, which is impossible since $(\mathcal{Y}^{\min})_s$ does not contain vanishing nodes by assumption, and its vertical (-2)-curves do not get contracted in $(\mathcal{Y}^{\mathrm{rst}})_s$ by construction. Finally, since we have $\mathcal{Y}^{\mathrm{rst}} \leq \mathcal{Y}^{\min}$, the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ contains no (-1)-line. Henceforth, the model $\mathcal{Y}^{\mathrm{rst}}$ is actually relatively stable.

Let us now prove the converse implication. Suppose that the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ exists. By definition, its specil fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ only contains persistent nodes; hence, by Proposition 3.7, the special fiber $(\mathcal{Y}^{\min})_s$ is obtained from $(\mathcal{Y}^{\mathrm{rst}})_s$ by replacing each of its node with an appropriate chain of horizontal (-2)-curves and persistent nodes and thus retains the property of not containing vanishing nodes. □

Proposition 3.12 establishes a criterion to determine whether Y has relatively stable reduction or not by looking at its minimal regular model. Hereafter we propose a refined version of such a criterion.

Lemma 3.13. Assume that $g(Y) \geq 1$. Then, given a regular or a semistable model $\mathcal{Y}$ of Y, no (-1)-line of $\mathcal{Y}_s$ can pass through a vanishing node of $\mathcal{Y}_s$.

Proof. Suppose that Q is a vanishing node, and let L be a (-1)-line of $\mathcal{Y}_s$ passing through it. Let $g \in G$ be an element stabilizing Q and flipping the two branches that pass through it. It is clear that gL will be another (-1)-line of $\mathcal{Y}_s$ passing through Q. Since $\mathcal{Y}_s$ contains two intersecting (-1)-lines, we have $g(C) = 0$ by Lemma 2.5, which contradicts our hypothesis. $\square$

Proposition 3.14. Assume that $g(Y) \geq 2$, or that $g(Y) = 1$ and $g(X) = 0$. The following are equivalent:

(a) Y has relatevely stable reduction;
(b) Y has semistable reduction, and the vanishing nodes of all models $\mathcal{Y}$ of Y acted upon by G all have even thickness;
(c) Y has semistable reduction, and the vanishing nodes of some semistable model $\mathcal{Y}$ of Y acted upon by G all have even thickness;
(d) Y has semistable reduction, and some semistable model $\mathcal{Y}$ of Y acted upon by G does not contain any vanishing node.

Proof. Let us prove (a) $\Longrightarrow$ (b). We will proceed by way of contradiction: we assume that there exists a model $\mathcal{Y}$ of Y acted upon by G which has a vanishing node Q of odd thickness t; we need to prove that $\mathcal{Y}^{\min}$ contains a vanishing node (see Proposition 3.12). We have that the minimal desingularization of $\mathcal{Y}$ is still semistable, and it also contains a vanishing node of odd thickness (see Proposition 3.7), hence we lose no generality if we assume that $\mathcal{Y}$ is regular. Let us further assume, without loss of generality, that $\mathcal{Y}$ is minimal (with respect to dominance) among the regular semistable models of Y acted upon by G and carrying a vanishing node Q. If $\mathcal{Y} = \mathcal{Y}^{\min}$, then we are done. If instead we have $\mathcal{Y} \gneq \mathcal{Y}^{\min}$, then $(\mathcal{Y})_s$ contains a G-orbit of (-1)-lines $\{gL : g \in G\}$; let $\mathcal{Y}_1$ be the semistable model that is obtained by contracting it. Since none of the gL can pass through the vanishing node Q by Lemma 3.13, the birational map $f : \mathcal{Y} \to \mathcal{Y}_1$ is an isomorphism above $Q_1 := f(Q)$; in particular, the node Q_1 is still a vanishing node of odd thickness of the model $\mathcal{Y}_1$, which, by construction, is again regular, semistable and acted upon by G; this violates the minimality of $\mathcal{Y}$.

The implication (b) $\Longrightarrow$ (c) is obvious; let us therefore prove (c) $\Longrightarrow$ (d). Let $\mathcal{Y}$ be some semistable model of Y whose vanishing nodes all have even thickness; by Proposition 3.7, the minimal desingularization $\mathcal{Y}'$ of $\mathcal{Y}$ will not contain a vanishing node, whence (d) follows.

Let us finally prove (d) $\Longrightarrow$ (a). If $\mathcal{Y}$ is a model acted upon by G which contains no vanishing nodes, its minimal desingularization has the same property by Proposition 3.7; hence we can assume without losing generality that $\mathcal{Y}$ is regular, and that it is moreover minimal (with respect to dominance) in the set of the regular semistable models of Y acted upon by G that do not contain vanishing nodes. If $\mathcal{Y} = \mathcal{Y}^{\min}$, we are done by Proposition 3.12. If instead we have $\mathcal{Y} \gneq \mathcal{Y}^{\min}$, then the special fiber $\mathcal{Y}_s$ contains a G-orbit of (-1)-lines $\{gL : g \in G\}$; if $\mathcal{Y}_1$ is the model we obtain by contracting them all, the birational map $\mathcal{Y} \to \mathcal{Y}_1$ is clearly an isomorphism above all nodes of $\mathcal{Y}_1$. The model $\mathcal{Y}_1$ will also not contain any vanishing node, and by construction it is still semistable and regular; this violates the minimality of $\mathcal{Y}$. $\square$

Corollary 3.15. Assume that $g(Y) \geq 2$, or that $g(Y) = 1$ and $g(X) = 0$. The curve Y always has relatively stable reduction over a large enough finite extension of R.

Proof. After possibly extending R, we may assume that Y has semistable reduction over R by Theorem 1.1. Let $\mathcal{Y}$ be any semistable model of Y acted upon by G. If the model $\mathcal{Y}$ does not contain a vanishing node of odd thickness, then the curve Y has relatively stable reduction over R by Proposition 3.14. Otherwise, let R' be any extension of R with even ramification index e. If we base-change $\mathcal{Y}$ to R', we still have a semistable model, and the thicknesses of the nodes all get multiplied by e. Hence, all nodes of $\mathcal{Y}_{R'}$ have even thickness, and Y has relatively stable reduction over R' thanks to Proposition 3.14. □

We end this subsection by pointing out a simple but important property of the relatively stable model.

Proposition 3.16. Suppose that Y has relatively stable reduction. If $W \in \mathrm{Irr}((\mathcal{X}^{(\mathrm{rst})})_s)$ is a smooth k-curve, then its inverse image in $(\mathcal{Y}^{\mathrm{rst}})_s$ also is. In particular, if all irreducible components of $(\mathcal{X}^{(\mathrm{rst})})_s$ are smooth k-curves, then the same is true of the irreducible components of $(\mathcal{X}^{(\mathrm{rst})})_s$.

Proof. Let Q be a node of $(\mathcal{Y}^{\mathrm{rst}})_s$ lying over a point $P \in W$. Since $\mathcal{Y}^{\mathrm{rst}}$ does not contain vanishing nodes by definition, we have that P is a node of $\mathcal{X}^{(\mathrm{rst})}$; moreover, since the component W is smooth at P, we have that P must connect W with another irreducible component $W' \in \mathrm{Irr}((\mathcal{X}^{(\mathrm{rst})})_s)$ distinct from W. We deduce from this that Q connects two distinct components V and V' of $(\mathcal{Y}^{\mathrm{rst}})_s$, one lying over W, and the other lying over W'. We conclude that, if V_1 and V_2 are two (possibly coinciding) irreducible components of $(\mathcal{Y}^{\mathrm{rst}})_s$ lying over W, they cannot be connected by a node; hence, the inverse image of W in $(\mathcal{Y}^{\mathrm{rst}})_s$ is a smooth k-curve. □

5. Finding the relatively stable model

Let us assume, for this subsection, that $g(Y) \geq 2$, or that $g(Y) = 1$ and $g(X) = 0$. The following result is the analog of Proposition 2.12 for the relatively stable model (instead of the minimal regular one).

Proposition 3.17. Assume that Y has relatively stable reduction, let $\mathcal{X}$ be any model of X, and let $\mathcal{Y}$ be the corresponding model of Y. Then, $\mathcal{X} \leq \mathcal{X}^{(\mathrm{rst})}$ if and only if, for all components V of $\mathcal{Y}_s$, we have $m(V) = 1$ and one of the following holds:

(i) $a(V) \geq 1$;
(ii) $a(V) = 0$ and $w(V) \geq 3$; or
(iii) $a(V) = 0$ and $\underline{w}(V) = (2)$.

Proof. Suppose first that we have $\mathcal{X} \leq \mathcal{X}^{(\mathrm{rst})}$. Given a component V of $\mathcal{Y}_s$, its strict transform V^{rst} in $\mathcal{Y}^{\mathrm{rst}}$ will have the same the same invariants m, a, w, and $\underline{w}$ as V by Lemmas 2.7 and 3.1. Now, since $(\mathcal{Y}^{\mathrm{rst}})_s$ is reduced, we have $m(V) = 1$. Since $(\mathcal{Y}^{\mathrm{rst}})_s$ does not contain (-1)-lines or horizontal (-2)-curves, we deduce from Lemma 2.8(c) and Definition 2.9 (applied to the model $\mathcal{Y}^{\mathrm{rst}}$) that one of the three conditions (i), (ii) and (iii) above must occur.

Now assume that $m(v) = 1$ and that either (i), (ii), or (iii) holds. By Proposition 2.12, we deduce that $\mathcal{Y} \leq \mathcal{Y}^{\min}$. Let V be any component of $\mathcal{C}_s$, and let $V^{\min}$ be the strict transform of V in $(\mathcal{Y}^{\min})_s$, which has the same invariants m, a, w, and $\underline{w}$ as V (Lemmas 2.7 and 3.1). Since we are excluding the case that $a(V) = 0$ and $\underline{w}(V) = (1, 1)$, the transform $V^{\min}$ is not a horizontal (-2)-curve (by Definition 2.9). Hence, all horizontal (-2)-curves of $(\mathcal{Y}^{\min})_s$ get contracted in $\mathcal{Y}_s$, which, in light of Proposition 3.12, is equivalent to saying that $\mathcal{Y} \leq \mathcal{Y}^{\mathrm{rst}}$. $\square$

Given any model $\mathcal{X}$ of X, Proposition 3.17 allows us to determine whether or not it is part of $\mathcal{X}^{(\mathrm{rst})}$. Meanwhile, the following proposition, allows us to determine the position of the components of $(\mathcal{X}^{(\mathrm{rst})})_s$ relative to the given model $\mathcal{X}$, under the assumption that $\mathcal{X} \leq \mathcal{X}^{(\min)}$.

Proposition 3.18. Assume that Y has relatively stable reduction, let $\mathcal{X}$ be any model of X, and let $\mathcal{Y}$ be the corresponding model of Y. Assume that $\mathcal{X} \leq \mathcal{X}^{(\min)}$. Then we have that $\mathrm{Ctr}(\mathcal{X}, \mathcal{X}^{(\mathrm{rst})})$ coincides with the set of those points P of $\mathcal{X}_s$ above which $\mathcal{Y}_s$ has non-nodal singularities or vanishing nodes.

Proof. Let us first remark that the minimal regular model $\mathcal{Y}^{\min}$ dominates both $\mathcal{Y}$ and $\mathcal{Y}^{\mathrm{rst}}$. Let Q be a point of $\mathcal{Y}_s$, and let P be its image in $\mathcal{X}_s$; as discussed in §3.1, we have $P \in \mathrm{Ctr}(\mathcal{X}, \mathcal{X}^{(\mathrm{rst})})$ if and only if $Q \in \mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$.

If Q is a non-singular point of $\mathcal{Y}_s$, then in particular it is a regular point of $\mathcal{Y}$, and hence the minimal desingularization morphism $\mathcal{Y}^{\min} \to \mathcal{Y}$ is an isomorphism above Q (i.e. no component of $(\mathcal{Y}^{\min})_s$ contracts to Q), and hence we have $Q \notin \mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$.

Suppose that Q is a persistent node of $\mathcal{Y}_s$. Then, its inverse image in $(\mathcal{Y}^{\min})_s$, by Proposition 3.7, consists of a chain of horizontal (-2)-curves, which will all be contracted in $(\mathcal{Y}^{\mathrm{rst}})_s$ (see Proposition 3.12). Hence, we have $Q \notin \mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$.

Suppose that Q is a vanishing node of $\mathcal{Y}_s$. Since Y has relatively stable reduction, it must have even thickness (see Proposition 3.14), and thus its inverse image in $(\mathcal{Y}^{\min})_s$ contains a vertical (-2)-curve by Proposition 3.7, which does not get contracted in $(\mathcal{Y}^{\mathrm{rst}})_s$ (see Proposition 3.12). Hence, we have $Q \in \mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$.

Suppose that Q is a non-nodal singularity of $\mathcal{Y}_s$; then the morphism $\mathcal{Y}^{\min} \to \mathcal{Y}$ cannot be an isomorphism above Q (because $\mathcal{Y}^{\min}$ is semistable), and the inverse image of Q in the semistable model $\mathcal{Y}^{\min}$ cannot only contain (-2)-curves; otherwise Q would be a node of $\mathcal{Y}_s$. Hence, there exists a component V of $(\mathcal{Y}^{\min})_s$ that contracts to Q and is not a (-2)-curve; by Proposition 3.12, it is clear that V does not get contracted in $(\mathcal{Y}^{\mathrm{rst}})_s$. Hence, we have $Q \in \mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$. $\square$

The following statement is the analog of Proposition 3.18 in the case $\mathcal{X} \not\leq \mathcal{X}^{(\min)}$.

Proposition 3.19. Assume that Y has relatively stable reduction, let $\mathcal{X}$ be any model of X, and let $\mathcal{Y}$ be the corresponding model of Y. Assume that $\mathcal{X} \not\leq \mathcal{X}^{(\min)}$, that $\mathcal{X}_s$ is irreducible and that $\mathcal{Y}_s$ is reduced. Then, $\mathcal{Y}_s$ has a unique singular point Q, which is a non-nodal singularity, and we have $\mathrm{Ctr}(\mathcal{X}, \mathcal{X}^{(\mathrm{rst})}) = \{f(Q)\}$, where $f : \mathcal{Y} \to \mathcal{X}$ is the cover map.

Proof. Since $\mathcal{X}$ has an irreducible special fiber, the components $\{V_i\}_{i=1}^N$ of $\mathcal{Y}_s$ are transitively permuted by G, and hence the invariants $m(V_i)$, $a(V_i)$ and $\underline{w}(V_i)$ do not depend on i. Since $\mathcal{Y}_s$ is reduced, we have $m(V_i) = 1$, and since $\mathcal{Y} \not\leq \mathcal{Y}^{\min}$, we have $a(V_i) = 0$ and $w(V_i) = 1$ by Proposition 2.12. This implies in particular that there is a unique singular point Q_i of $\mathcal{Y}_s$ that lies on V_i; moreover, since $\mathcal{Y}_s$ is connected, the Q_i's must all coincide.

We conclude that $\mathcal{Y}_s$ contains a unique singular point Q. If Q were a node, then $\mathcal{Y}_s$ would be semistable, and its vertical components would all be (-1)-lines by Lemma 2.8(c), which is impossible since $g(Y) \geq 1$. Hence, Q is a non-nodal singularity of $\mathcal{Y}_s$. Let $\mathcal{Y}'$ be the minimal desingularization of $\mathcal{Y}$; we have dominance relations $\mathcal{Y} \leq \mathcal{Y}' \gneq \mathcal{Y}^{\min} \geq \mathcal{Y}^{\mathrm{rst}}$. Since $\mathcal{Y} \not\leq \mathcal{Y}^{\min}$, we have $\mathcal{Y} \not\leq \mathcal{Y}^{\mathrm{rst}}$, which implies that $\mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}}) \neq \varnothing$. Since the desingularization $\mathcal{Y}' \to \mathcal{Y}$ is necessarily an isomorphism above the smooth points of $\mathcal{Y}_s$, the set $\mathrm{Ctr}(\mathcal{Y}, \mathcal{Y}^{\mathrm{rst}})$ cannot contain any point of $\mathcal{Y}_s$ other then Q. □

When $\mathcal{X}$ is a smooth model, the (other) components of $\mathcal{X}^{(\mathrm{rst})}$ are contracted precisely to the points of $\mathcal{X}_s$ above which $\mathcal{Y}_s$ is singular, as the following corollary emphasizes.

Corollary 3.20. Assume that Y has relatively stable reduction, let $\mathcal{X}$ be any model of X, and let $\mathcal{Y}$ be the corresponding model of Y. Assume that $\mathcal{Y}$ has reduced special fiber. Then we have $\mathrm{Ctr}(\mathcal{X}, \mathcal{X}^{(\mathrm{rst})}) = f(\mathrm{Sing}(\mathcal{Y}_s))$, where $f : \mathcal{Y} \to \mathcal{X}$ is the covering map and $\mathrm{Sing}(\mathcal{Y}_s)$ is the set of all singular points of $\mathcal{Y}_s$.

Proof. This is immediate from Propositions 3.18 and 3.19. □

CHAPTER 4

Models of hyperelliptic curves

In this section, we specialize to the case in which the Galois cover $Y \to X$ that we considered in §3 is the degree-2 map from a hyperelliptic curve Y of genus $g \geq 1$ to the projective line X (see §1.5.5.2). Our main aim, for this and the following sections, is computing the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ of Y that we defined in §3.4. After recalling some basic general facts about hyperelliptic curves in §4.1, we describe the semistable models of the line X in §4.2: we will see how smooth models of the line correspond to discs $D \subseteq \bar{K}$, while the semistable ones correspond to certain finite collections $\mathfrak{D}$ of discs. After introducing the notion of a *part-square decomposition* of a polynomial in §4.3, we exploit it in §4.4 to describe the model of the hyperelliptic curve Y corresponding to a given smooth model of the line X (i.e., to a given disc D). The special fiber of such models of Y will be more thoroughly studied in §4.5 and §4.6, and each of the two subsections provides a criterion to decide whether a disc D belongs to the collection $\mathfrak{D}$ that defines the semistable model of the line $\mathcal{X}^{(\mathrm{rst})}$ from which $\mathcal{Y}^{\mathrm{rst}}$ comes (Theorem 4.32 for the separable case, and Proposition 4.35 for the inseparable case).

1. Equations for hyperelliptic curves

We let F be any field and write $X := \mathbb{P}^1_F$ for the projective line over F. In this subsection, we review basic facts and definitions relating to hyperelliptic curves over F which can be found in [11, §7.4.3].

Definition 4.1. A *hyperelliptic curve* over F is a smooth curve Y/F of genus $g \geq 1$ along with a separable (branched) covering morphism $Y \to X$ of degree 2, which we call the *hyperelliptic map*.

It is possible through repeated applications of the Riemann-Roch Theorem to show the well-known fact that the affine chart $x \neq \infty$ of any hyperelliptic curve Y/F can be described by an equation of the form

$$y^2 + q(x)y = r(x), \tag{10}$$

where $\deg(q) \leq g+1$ and $\deg(r) \leq 2g+2$ and the hyperelliptic map is given by the coordinate $x : Y \to X$. If F has characteristic different from 2, a suitable change of the coordinate y allows us to convert this equation into the simpler form $y^2 = r(x) + \frac{1}{4}q^2(x)$, from which it is clear that the smoothness condition implies that the polynomial $f(x) := r(x) + \frac{1}{4}q^2(x)$ must be separable. Over the complementary affine chart of $\mathbb{P}^1_F$ where $x \neq 0$, the hyperelliptic curve Y can be described by the equation

$$\check{y}^2 + \check{q}(\check{x})\check{y} = \check{r}(\check{x}), \tag{11}$$

where $\check{x} = x^{-1}$, $\check{y} = x^{-(g+1)}y$, $\check{q}(\check{x}) = x^{-(g+1)}q(x)$, and $\check{r}(\check{x}) = x^{-(2g+2)}r(x)$. Note that the polynomial $\check{q}(z) \in F[z]$ (resp. $\check{r}(z) \in F[z]$) differs from the polynomial $q(z) \in F[z]$ (resp.

$r(z) \in F[z]$) only in that each power z^i which appears in the polynomial is replaced by z^{g+1-i} (resp. z^{2g+2-i}) while the coefficients remain the same.

If F has characteristic different from 2 (i.e. if we consider *tame* hyperelliptic curves), the Riemann-Hurwitz formula ensures that the ramification locus of $Y_{\bar{F}} \to X_{\bar{F}}$ consists of $2g+2$ points of $Y_{\bar{F}}$, lying over $2g+2$ distinct branch points of $X_{\overline{F}}$. In fact, the branch locus determines a hyperelliptic curve almost completely, as we see from the following proposition.

Proposition 4.2. Given a field F of characteristic different from 2, and letting X be the projective line $\mathbb{P}^1_F$, the following data are equivalent:

(i) a hyperelliptic curve Y of genus g having rational branch locus, endowed with a distinguished hyperelliptic map $Y \to X$;

(ii) a separable polynomial $f(x) \in F[x]$ of degree $2g+1$ or $2g+2$ all of whose roots lie in F, modulo multiplication by a scalar in $(F^\times)^2$; and

(iii) a cardinality-$(2g+2)$ subset $\mathcal{B} \subset X(F)$ together with an element $c \in F^\times/(F^\times)^2$.

Moreover, in (ii) above, the polynomial f will have degree $2g+1$ (resp. $2g+2$) if in the context of (iii) above the coordinate of X is chosen such that ∞ is (resp. is not) an element of $\mathcal{B}$.

Proof. We construct the above equivalences as follows. Given a hyperelliptic curve Y as in (i), we denote the distinguished hyperelliptic map by $x : Y \to X$. Clearly, the morphism x can be viewed as an element of the function field $F(Y)$; in fact, as the hyperelliptic map is not constant, we must have $F(X) = F(x) \hookrightarrow F(Y)$. Since the hyperelliptic map has degree 2, the extension $F(Y) \supset F(X)$ must be generated by a single element $y \in F(Y) \smallsetminus F(X)$ with $y^2 \in F(X)$; after multiplying y by a suitable polynomial in x, we may assume that $f(x) := y^2 \in F[x]$. Then it is straightforward to see that the affine open subset of Y given by the inverse image of $\mathbb{A}^1_F$ under the hyperelliptic map is described by the equation $y^2 = f(x)$. The roots of f clearly coincide with the points on $\mathbb{A}^1_{\bar{F}}$ over which the hyperelliptic map is ramified; meanwhile, one sees by applying the Riemann-Hurwitz formula that the map $Y \to X$ must have exactly $2g+2$ ramification points. Therefore, the polynomial f has $2g+1$ (resp. $2g+2$) roots all lying in F if $\infty \in X(F)$ is (resp. is not) a ramification point. This polynomial f (modulo multiplication by elements in $(F^\times)^2$) gives us the data in (ii).

Given the data in (ii), let $\mathcal{R} \subset X(F) \smallsetminus \{\infty\}$ be the subset of roots of f and let $c \in F^\times/(F^\times)^2$ be the leading coefficient of f modulo squares of elements in $F^\times$. Setting $\mathcal{B} = \mathcal{R}$ (resp. $\mathcal{B} = \mathcal{R} \cup \{\infty\}$) if the degree of f is $2g+1$ (resp. $2g+2$), we have that the set $\mathcal{B}$ has cardinality $2g+2$ and we get the data of (iii).

Finally, given a cardinality-$(2g+2)$ subset $\mathcal{B} \subset X(F)$ and a scalar $c \in F^\times/(F^\times)^2$, write $\tilde{c} \in F^\times$ for a representative of c and let Y/F be the smooth completion of the affine curve over F described by the equation

$$y^2 = f(x) := \tilde{c} \prod_{a \in \mathcal{B} \smallsetminus \{\infty\}} (x-a). \tag{12}$$

Then it is easy to check that Y satisfies the criteria given in (i), with the hyperelliptic map being given by the function $x \in F(Y)$. If a different representative $\tilde{c}' \in F^\times$ is chosen for c in order to define f, then we must have $\tilde{c}' = \gamma^2 \tilde{c}$ for some $\gamma \in F^\times$, and replacing the coordinate y by γy gives us the same equation (12) and therefore the same curve Y, so the data in (i) is uniquely determined by (iii). □

Remark 4.3. Suppose that in the context of the above proposition, none of branch points of $Y \to X$ is ∞. We may then find an isomorphic hyperelliptic curve over F whose branch points over X include the point $\infty \in X(F)$ by applying an automorphism of the projective line X which moves one of the branch points to ∞. More precisely, if a_0 is the x-coordinate of a branch point of $Y \to X$ that does not coincide with ∞, we perform the substitution $(x, y) \mapsto ((x - a_0)^{-1}, \check{x}^{g+1}\check{y})$ and get a curve (isomorphic over F to our original one) ramified over $\infty \in X(F)$ defined by an equation of the form $y^2 = f(x)$, where $f(x) \in F[x]$ is a polynomial of degree $2g + 1$.

From now on, we assume that F is the discretely-valued field K satisfying the conditions given in §1.5. In light of the remark above, up to possibly replacing K with a finite extension (so that at least one of the branch points of the cover $Y \to X$ is rational), we can and will make the following assumption throughout the rest of the work.

Hypothesis 4.4. The hyperelliptic curve Y is defined over K by the equation $y^2 = f(x)$, where x is the standard coordinate of $X = \mathbb{P}^1_K$, and $f(x) \in K[x]$ is a polynomial of (odd) degree $2g + 1$, where g is the genus of Y.

Proposition 4.2 allows us to treat the hyperelliptic curve Y essentially as a marked line. This is peculiar to the hyperelliptic case: if we were to deal with a tame Galois covering of the line of degree greater than 2, the same branch locus would in general be shared by multiple branched coverings corresponding to various possible monodromy actions.

Our aim will be constructing semistable models of a given hyperelliptic curve $Y \to X$ by normalizing some carefully chosen semistable models of the line X in the quadratic extension $K(X) \subseteq K(Y)$. We will start by analyzing smooth and semistable models of the line X in the next subsection; in the subsequent ones, we will turn our attention to the corresponding models of Y.

2. Models of the projective line

As before, let $X := \mathbb{P}^1_K$ be the projective line, and let x denote its standard coordinate. Given $\alpha \in \bar{K}$ and $\beta \in \bar{K}^\times$, one can define a smooth model $\mathcal{X}_{\alpha,\beta}$ of X over the ring of integers R' of $K' := K(\alpha, \beta) \subseteq \bar{K})$ by declaring $\mathcal{X}_{\alpha,\beta} := \mathbb{P}^1_{R'}$, with coordinate $x_{\alpha,\beta} := \beta^{-1}(x - \alpha)$, as an R'-scheme, and identifying the generic fiber $\mathcal{X}_\eta$ with X via the linear transformation $x_{\alpha,\beta} = \beta^{-1}(x - \alpha)$. If (α_1, β_1) and (α_2, β_2) are such that $v(\alpha_1 - \alpha_2) \leq v(\beta_1) = v(\beta_2)$, then $\mathcal{X}_{\alpha_1,\beta_1}$ and $\mathcal{X}_{\alpha_2,\beta_2}$ are isomorphic as models of X, the isomorphism being given by the change of variable $x_{\alpha_2,\beta_2} = u x_{\alpha_1,\beta_1} + \delta$, where u is the unit $\beta_1(\beta_2)^{-1}$ and δ is the integral element $\beta_2^{-1}(\alpha_1 - \alpha_2)$. In other words, the model $\mathcal{X}_{\alpha,\beta}$ only depends, up to isomorphism, on the disc $D = D_{\alpha,b} := \{x \in \bar{K} : v(x - \alpha) \leq b\}$ of center α and depth $b := v(\beta)$; for this reason, we will often denote $\mathcal{X}_{\alpha,\beta}$ by $\mathcal{X}_D$.

Proposition 4.5. The construction $D \mapsto \mathcal{X}_D$ described above defines a bijection between the discs of $\bar{K}$ and the smooth models of X defined over finite extensions of R considered up to isomorphism (two models $\mathcal{X}_1/R'_1$ and $\mathcal{X}_2/R'_2$ are considered isomorphic if they become so over some common finite extension $R'' \supseteq R'_1, R'_2$).

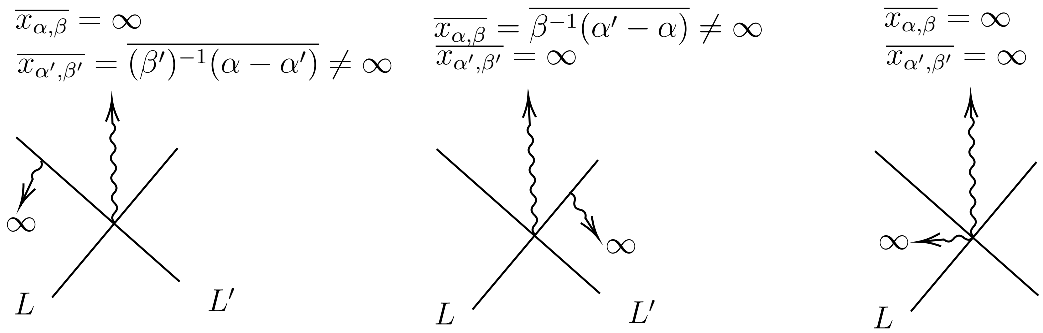

Case (a): $D \subsetneq D'$ Case (b): $D' \subsetneq D$ Case (c): $D \cap D' = \emptyset$

FIGURE 1. The special fiber of the minimal model $\mathcal{X}_{\{D,D'\}}$ dominating $\mathcal{X}_D$ and $\mathcal{X}'_D$. Here, L and L' are the lines corresponding to the discs D and D' respectively.

PROOF. Given a smooth model of the line $\mathcal{X}$ over the ring of integers R' of some finite extension K'/K, one may prove that there exists an isomorphism of R'-schemes $\mathcal{X} \cong \mathbb{P}^1_{R'}$ (see [11, Exercise 8.3.5]), which immediately implies that $\mathcal{X}$ is isomorphic, as a model, to $\mathcal{X}_D$ for some uniquely determined disc $D = D_{\alpha,b}$ with $\alpha \in K'$ and $b \in v((K')^\times)$. $\square$

Given two discs $D = D_{\alpha,b}$ and $D' = D_{\alpha',b'}$, we want to compare the smooth models $\mathcal{X}_D$ and $\mathcal{X}_{D'}$: using the notation introduced in §2.1.1.3, we have the following proposition, which can be verified by an immediate computation.

Proposition 4.6. With the notation above, assume $D \neq D'$, and let $P \in (\mathcal{X}_D)_s(k)$ and $P' \in (\mathcal{X}_{D'})_s(k)$ be the points such that $\mathrm{Ctr}(\mathcal{X}_D, \mathcal{X}_{D'}) = \{P\}$ and $\mathrm{Ctr}(\mathcal{X}_{D'}, \mathcal{X}_D) = \{P'\}$. Then there are the following three possibilities (illustrated in Figure 1):

(a) when $D \subsetneq D'$, P is the point $\overline{x_{\alpha,\beta}} = \infty$ and P' is the point $\overline{x_{\alpha',\beta'}} = \overline{(\beta')^{-1}(\alpha - \alpha')} \neq \infty$;
(b) when $D' \subsetneq D$, P is the point $\overline{x_{\alpha,\beta}} = \overline{\beta^{-1}(\alpha' - \alpha)} \neq \infty$, and P' is the point $\overline{x_{\alpha',\beta'}} = \infty$; or
(c) when $D \cap D' = \varnothing$, P is the point $\overline{x_{\alpha,\beta}} = \infty$ and P' is the point $\overline{x_{\alpha',\beta'}} = \infty$.

If $\mathfrak{D} = \{D_1, \ldots, D_n\}$ is a non-empty, finite collection of discs of $\bar{K}$, one can form a corresponding model $\mathcal{X}_{\mathfrak{D}}$, which is defined as the minimal model dominating all the smooth models $\{\mathcal{X}_D : D \in \mathfrak{D}\}$. Its special fiber is a reduced k-curve of arithmetic genus $p_a((\mathcal{X}_{\mathfrak{D}})_s) = g(X) = 0$, i.e. it consists of n lines $L_1, \ldots, L_n$ corresponding to the discs D_i's meeting each other at ordinary multiple points, without forming loops.

Proposition 4.7. The construction $\mathfrak{D} \mapsto \mathcal{X}_{\mathfrak{D}}$ described above defines a bijection between the finite non-empty collection of discs $\mathfrak{D}$ of $\bar{K}$ and the models of X having reduced special fiber defined over finite extensions of R, considered up to isomorphism (two models $\mathcal{X}_1/R'_1$ and $\mathcal{X}_2/R'_2$ are considered isomorphic if they become so over some common finite extension $R'' \supseteq R'_1, R'_2$).

PROOF. Suppose that $\mathcal{X}$ is a model of the line X with reduced special fiber, and let $\{L_1, \ldots, L_n\}$ be the components of its special fiber $\mathcal{X}_s$. For each i, let $\mathcal{X}_i$ be the model of obtained from $\mathcal{X}$ by contracting all lines in $\mathrm{Irr}(\mathcal{X}_s)$, except for L_i it is easy to prove that $\mathcal{X}_i$ is smooth (see, for example, [11, Exercise 8.3.5]), so that, by Proposition 4.5, $\mathcal{X}_i = \mathcal{X}_{D_i}$ for some uniquely determined disc $D_i \subset \bar{K}$. Now the model $\mathcal{X}$ can be described as the minimal model dominating all the $\mathcal{X}_i$'s, i.e. we have $\mathcal{X} \cong \mathcal{X}_{\mathfrak{D}}$ with $\mathfrak{D} = \{D_1, \ldots, D_n\}$. $\square$

Proposition 4.8. Suppose that $\mathcal{X}/R'$ is a semistable model of the line X for some finite extension R'/R, such that there are discs $D_{\alpha,b} \subsetneq D_{\alpha',b'} \subset \bar{K}$ corresponding to two intersecting components of $(\mathcal{X})_s$. Then the thickness of the node where they intersect is given by the formula $(b'-b)/v(\pi)$, where $\pi \in \bar{K}$ is a uniformizer of R'.

Proof. We can clearly replace the center α' with $\alpha \in D_{\alpha,b} \subsetneq D_{\alpha',b'}$; now choosing $\beta, \beta' \in \bar{K}^\times$ be scalars such that $v(\beta) = b$ and $v(\beta') = b'$. Then, with the notation above, we have coordinates $x_{\alpha,\beta}$ and $x_{\alpha,\beta'}$ corresponding to each of these components of $(\mathcal{X})_s$, and these coordinates are related by the equation $x_{\alpha,\beta} = \beta'\beta^{-1}x_{\alpha,\beta'}$. Locally around the point of intersection, a defining equation is $x_{\alpha,\beta}x^\vee_{\alpha,\beta'} = \beta'\beta^{-1}$, where $x^\vee_{\alpha,\beta'} = x^{-1}_{\alpha,\beta'}$, and so the thickness by definition is equal to $v(\beta'\beta^{-1})/v(\pi) = (b'-b)/v(\pi)$. $\square$

Proposition 4.7 certainly implies that a semistable model of the line always has the form $\mathcal{X}_{\mathfrak{D}}$ for some finite non-empty family of discs $\mathfrak{D}$; however, it is not always true that, given a collection of discs $\mathfrak{D}$, the corresponding model $\mathcal{X}_{\mathfrak{D}}$ of the line is semistable. In fact, its special fiber $(\mathcal{X}_{\mathfrak{D}})_s$ is always a k-curve with at worst ordinary singularities, but it is possible that more than two lines $L_i \in \mathrm{Irr}((\mathcal{X}_{\mathfrak{D}})_s)$ intersect at the same ordinary multiple point, violating semistability. However, it is not difficult to give a combinatorial necessary and sufficient condition for a collection $\mathfrak{D}$ of discs to give rise to a semistable model.

Proposition 4.9. The model $\mathcal{X}_{\mathfrak{D}}$ of X corresponding to a finite non-empty collection of discs $\mathfrak{D}$ is semistable if and only if it satisfies the following property: if three discs $D_1, D_2, D_3 \in \mathfrak{D}$ satisfy any of the three conditions

(a) $D_1, D_2, D_3 \in \mathfrak{D}$ are mutually disjoint, and any disc in $\bar{K}$ containing two of them also contains the third one;
(b) $D_1, D_2, D_3 \in \mathfrak{D}$ are mutually disjoint, and there exists a disc in $\bar{K}$ containing D_1 and D_2 that is disjoint from D_3; or
(c) $D_3 \supsetneq D_1 \cup D_2$, and D_3 is not minimal among the discs of $\bar{K}$ satisfying this property,

then, letting D be the minimal disc of $\bar{K}$ containing both D_1 and D_2, we have $D \in \mathfrak{D}$.

The proof of the proposition relies on the following elementary lemma.

Lemma 4.10. We have the following.

(a) Given three discs $D_1, D_2, D_3 \subset \bar{K}$, some permutation of them satisfies the assumptions (a), (b) or (c) of Proposition 4.9 if and only if $(\mathcal{X}_{\{D_1,D_2,D_3\}})_s$ consists of three lines L_1, L_2 and L_3 meeting at an ordinary triple point.
(b) Given three discs D_1, D_2 and D_3 satisfying the assumptions (a), (b) or (c) of Proposition 4.9, and letting D be the minimal disc containing D_1 and D_2, we have that the three lines L_1, L_2 and L_3 corresponding to D_1, D_2, and D_3 do not intersect each other in the special fiber of $\mathcal{X}_{\{D_1,D_2,D_3,D\}}$, and they intersect the line L corresponding to the disc D at three distinct points (see Figure 2).

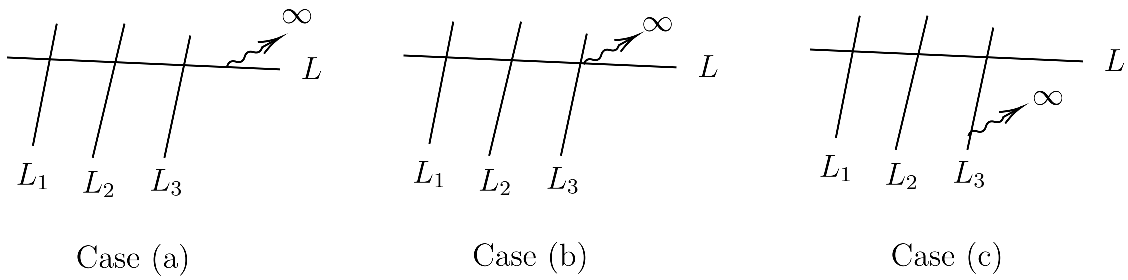

FIGURE 2. When three discs D_1, D_2, D_3 satisfy the assumptions of points (a), (b) or (c) of Proposition 4.9. and D is the minimal disc containing D_1 and D_2, then the special fiber of $\mathcal{X}_{\{D_1,D_2,D_3,D\}}$, has the shape depicted above (the converse is actually also true). In the picture, L_i is the line corresponding to the disc D_i, and L is the line corresponding to the disc D.

PROOF. The lemma can be proved by means of straightforward computations, which we omit. □

PROOF OF PROPOSITION 4.9. Both implications will be proved by way of contradiction.

Assume that $\mathcal{X}_{\mathfrak{D}}$ is not semistable, so that there exists an ordinary singular point $P \in (\mathcal{X}_{\mathfrak{D}})_s$ through which three distinct lines $L_1, L_2, L_3 \in \mathrm{Irr}((\mathcal{X}_{\mathfrak{D}})_s)$ pass; letting D_1, D_2 and D_3 the corresponding three discs, Lemma 4.10(a) ensures that they satisfy (possibly after performing a permutation) condition (a), (b) or (c) of Proposition 4.9. Now if D is the minimal disc containing D_1 and D_2, then we certainly have $D \notin \mathfrak{D}$: otherwise, we would have $\mathcal{X}_{\mathfrak{D}} \geq \mathcal{X}_{\{D_1,D_2,D_3,D\}}$, and this would prevent L_1, L_2 and L_3 from intersecting each other in $\mathcal{X}_{\mathfrak{D}}$ by Lemma 4.10(b).

Conversely, assume that $D_1, D_2, D_3 \in \mathfrak{D}$ are three discs satisfying either condition (a), (b) or (c) of Proposition 4.9, and such that $D \notin \mathfrak{D}$, where D is the minimal disc containing D_1 and D_2. Let P be the point of $(\mathcal{X}_{\mathfrak{D}})_s$ such that $\mathrm{Ctr}(\mathcal{X}_{\mathfrak{D}}, \mathcal{X}_D) = \{P\}$; in other words, P is the point of $(\mathcal{X}_{\mathfrak{D}})_s$ to which the unique line L of which the special fiber of $\mathcal{X}_D$ consists is contracted. We observe that, in the model $\mathcal{X}_{\{D_1,D_2,D_3,D\}}$, the line L corresponding to the disc D intersects the rest of the special fiber at more than 2 points (this follows from Lemma 4.10(b)); the same will consequently also be true in the model $\mathcal{X}_{\mathfrak{D}\cup\{D\}} \geq \mathcal{X}_{\{D_1,D_2,D_3,D\}}$. Hence, at least 3 lines will pass through the point $P \in (\mathcal{X}_{\mathfrak{D}})_s$ to which the line $L \in \mathrm{Irr}((\mathcal{X}_{\mathfrak{D}\cup\{D\}})_s)$ gets contracted, which implies that $\mathcal{X}_{\mathfrak{D}}$ is not semistable. □

The following way of rephrasing the conditions (a), (b) and (c) of Proposition 4.9 will also be useful later.

Proposition 4.11. Given a collection of discs $\mathfrak{D}$ and a disc D of $\bar{K}$, the following are equivalent:

(i) there exist discs $D_1, D_2, D_3 \in \mathfrak{D}$ which satisfy conditions (a), (b) or (c) of Proposition 4.9 and the disc D is the minimal disc of $\bar{K}$ containing D_1 and D_2;
(ii) we have $|\,\mathrm{Ctr}(\mathcal{X}_D, \mathcal{X}_{\mathfrak{D}})| \geq 3$.

PROOF. First assume that (i) holds. It follows from Lemma 4.10(b) that $|\operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\{D_1,D_2,D_3\}})| = 3$, which implies (ii), since we clearly have $\operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\{D_1,D_2,D_3\}}) \subseteq \operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\mathfrak{D}})$. Now assume that (ii) holds; from $|\operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\mathfrak{D}})| \geq 3$ one deduces that there exist discs D_1, D_2, and $D_3 \in \mathfrak{D}$ such that $\operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\{D_1,D_2,D_3\}}) = \bigcup_{i=1}^3 \operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{D_i})$ consists of three distinct points of $(\mathcal{X}_D)_s$; now (i) follows from straightforward calculations, taking into account Proposition 4.6. □

Given any non-empty collection of discs $\mathfrak{D}$, we can complete it to a family $\mathfrak{D}^{\mathrm{ss}}$ of discs corresponding to a semistable model: it is enough that, for every three discs $D_1, D_2, D_3 \in \mathfrak{D}$ satisfying the conditions (a), (b) or (c) of Proposition 4.9, the minimal disc containing D_1 and D_2 is added to $\mathfrak{D}$. It is not difficult to see that the resulting family of discs $\mathfrak{D}^{\mathrm{ss}} \supseteq \mathfrak{D}$ satisfies the hypothesis of Proposition 4.9 and consequently corresponds to the minimal semistable model $\mathcal{X}_{\mathfrak{D}^{\mathrm{ss}}}$ of the line X that dominates $\mathcal{X}_{\mathfrak{D}}$.

Remark 4.12. Suppose a non-empty collection of discs $\mathfrak{D}$ corresponding to a semistable model $\mathcal{X}_{\mathfrak{D}}$ is given, and let $\mathfrak{D}' \subseteq \mathfrak{D}$ be a non-empty subfamily. Suppose that, for all $D \in \mathfrak{D} \smallsetminus \mathfrak{D}'$, there exists three discs $D_1, D_2, D_3 \in \mathfrak{D}$ satisfying the conditions (a), (b) or (c) of Proposition 4.9, and such that D is the minimal disc of $\bar{K}$ containing D_1 and D_2 – by Proposition 4.11, this condition can be equivalently expressed by saying that, for all $D \in \mathfrak{D} \smallsetminus \mathfrak{D}'$, the set $\operatorname{Ctr}(\mathcal{X}_D, \mathcal{X}_{\mathfrak{D}})$ consists of three or more points. Then, $\mathfrak{D}$ can be reconstructed from $\mathfrak{D}'$ by applying the completion procedure described above, i.e. we have $\mathfrak{D} = (\mathfrak{D}')^{\mathrm{ss}}$.

3. Part-square decompositions

We begin this subsection by defining a *part-square decomposition*, and then we study part-square decompositions with certain properties.

Definition 4.13. Given a nonzero polynomial $h(x) \in \bar{K}[z]$, a *part-square decomposition* of h is a way of writing $h = q^2 + \rho$ for some $q(x), \rho(x) \in \bar{K}[x]$, with $\deg(q) \leq \lceil \deg(h)/2 \rceil$.

Remark 4.14. The definition forces $\deg(\rho) \leq \deg(h)$ when h has even degree and $\deg(\rho) \leq \deg(h) + 1$ when h has odd degree. The definition allows q to be equal to zero.

Given a part-square decomposition $h = q^2 + \rho$, we define the rational number $t_{q,\rho} := v(\rho) - v(h) \in \mathbb{Q} \cup \{+\infty\}$.

Definition 4.15. We define the following properties of a part-square decomposition $h = q^2 + \rho$.

(a) The decomposition is said to be *good* either if we have $t_{q,\rho} \geq 2v(2)$ or if we have $t_{q,\rho} < 2v(2)$ and there is no decomposition $h = \tilde{q}^2 + \tilde{\rho}$ such that $t_{\tilde{q},\tilde{\rho}} > t_{q,\rho}$.

(b) The decomposition is said to be *totally odd* if ρ only consists of odd-degree terms.

Remark 4.16. The trivial part-square decomposition $h = 0^2 + h$ has $t_{0,h} = 0$; this immediately implies that all good decompositions $h = q^2 + \rho$ satisfy $t_{q,\rho} \geq 0$. When $p \neq 2$, the converse also holds because we have $2v(2) = 0$.

Remark 4.17. If $h = q^2 + \rho = (q')^2 + \rho'$ are two good part-square decompositions for the same nonzero polynomial h, then we have $\min\{t_{q,\rho}, 2v(2)\} = \min\{t_{q',\rho'}, 2v(2)\}$ directly from Definition 4.15.

Proposition 4.18. Let $h = q^2 + \rho$ be a part-square decomposition satisfying $t_{q,\rho} < 2v(2)$. Then we have the following.

(a) The decomposition $h = q^2 + \rho$ is good if and only if the normalized reduction of ρ is not the square of a polynomial with coefficients in k.

(b) Suppose that the decomposition $h = q^2 + \rho$ is good and that $h = \tilde{q}^2 + \tilde{\rho}$ is another good decomposition. Then given any normalized reductions of ρ and $\tilde{\rho}$ respectively, the same odd degrees appear among terms in these normalized reductions, and their derivatives are equal up to scaling.

Proof. We begin by proving part (a). If $t_{q,\rho} < 0$, the decomposition is not good (see Remark 4.16), and any normalized reduction of ρ is a square, since it is a scalar multiple of a normalized reduction of q^2. We now have to prove the two implications when $t_{q,\rho} \geq 0$.

Suppose that $h = q^2 + \rho$ satisfies $0 \leq t_{q,\rho} < 2v(2)$ but is not good, so that another decomposition $h = \tilde{q}^2(x) + \tilde{\rho}(x)$ with $t_{\tilde{q},\tilde{\rho}} > t_{q,\rho}$ can be found. Let us now consider $q + \tilde{q}$ and $q - \tilde{q}$: their product has valuation $v(q^2 - \tilde{q}^2) = v(\tilde{\rho} - \rho) = v(\rho)$, while their difference has valuation

$$v(2\tilde{q}) = v(2) + \frac{1}{2}v(h - \tilde{\rho}) \geq v(2) + \frac{1}{2}v(h) > \frac{1}{2}v(\rho). \tag{13}$$

From this, it is immediate to deduce that they must both have valuation equal to $\frac{1}{2}v(\rho)$. We may now write

$$\rho = \tilde{\rho} + 2\tilde{q}(\tilde{q} - q) - (\tilde{q} - q)^2. \tag{14}$$

But we observe that that the first two summands both have valuation $> v(\rho)$. This implies that the normalized reduction of ρ is a square.

Conversely, suppose that the decomposition satisfies $0 \leq t_{q,\rho} < 2v(2)$ and that the normalized reduction of ρ is a square; this is clearly equivalent to saying that we can form a part-square decomposition $\rho = q_1^2 + \rho_1$ of the polynomial ρ that satisfies $t_{q_1,\rho_1} > 0$; hence, we have $v(q_1) = \frac{1}{2}v(\rho)$ and $v(\rho_1) > v(\rho)$.

Let us now consider the part-square decomposition $h = \tilde{q}^2 + \tilde{\rho}$, where $\tilde{q} := q + q_1$ and $\tilde{\rho} = \rho_1 - 2qq_1$. Notice that the assumption $t_{q,\rho} \geq 0$ implies that $v(q) \geq v(h)/2$; we therefore have $v(2qq_1) \geq v(2) + \frac{1}{2}v(h) + \frac{1}{2}v(\rho) > v(\rho)$. We conclude that $v(\tilde{\rho}) = v(\rho_1 - 2qq_1) > v(\rho)$, i.e. $t_{\tilde{q},\tilde{\rho}} > t_{q,\rho}$; therefore, the original part-square decomposition $h = q^2 + \rho$ was not good. Thus, both directions of part (a) are proved.

We now turn to part (b) and assume that $h = q^2 + \rho = \tilde{q}^2 + \tilde{\rho}$ are both good decompositions. Then both (13) and (14) are still valid, and the fact that $v(\rho) = v(\tilde{\rho})$ implies that $v(2\tilde{q}(\tilde{q} - q) - (q - \tilde{q})^2) \geq v(\rho)$. Then if $v(q - \tilde{q}) < \frac{1}{2}v(\rho)$, from (14) we must have $v(2\tilde{q}(\tilde{q} - q)) = v((q - \tilde{q})^2) < v(\rho)$, which contradicts (13). We therefore have $v(\tilde{q} - q) \geq \frac{1}{2}v(\rho)$, from which $v(2\tilde{q}(\tilde{q} - q)) > \frac{1}{2}v(\rho)$ follows from (13).

Let $\gamma \in \bar{K}$ be a scalar with $v(\gamma) = v(\rho) = v(\tilde{\rho})$. If $v(\tilde{q} - q) > \frac{1}{2}v(\rho)$, then (14) shows that $v(\tilde{\rho} - \rho) > v(\rho)$ and so $\gamma^{-1}(\tilde{\rho} - \rho)$ has positive valuation; therefore, the reductions of $\gamma^{-1}\rho$ and $\gamma^{-1}\tilde{\rho}$ are equal, and we are done. If $v(\tilde{q} - q) = \frac{1}{2}v(\rho)$, then $\gamma^{-1}(\tilde{\rho} - \rho)$ reduces to a square (namely a normalized reduction of $\tilde{q} - q$ squared); the square of a polynomial in $k[x]$ has only even-degree terms, and its derivative vanishes, which shows that the reductions of $\gamma^{-1}\rho$ and $\gamma^{-1}\tilde{\rho}$ have the same odd degrees appearing and have the same derivative. Thus again we are done, and part (b) is proved. $\square$

Corollary 4.19. Every totally odd part-square decomposition of a polynomial is good.

Proof. Suppose that the decomposition $h = q^2 + \rho$ is totally odd. If $t_{q,\rho} \geq 2v(2)$, then we are already done, so assume that $t_{q,\rho} < 2v(2)$. Then since ρ consists only of odd-degree terms, the same is true of any normalized reduction of ρ, which consequently cannot be the square of any polynomial in $k[z]$. Then Proposition 4.18 implies that the decomposition is good. $\square$

We now want to show that a good part-square decomposition of a polynomial always exists, for which, thanks to Corollary 4.19, it suffices to show that a polynomial always has a totally odd part-square decomposition.

Proposition 4.20. Given a nonzero polynomial $h(z) \in \bar{K}[z]$, there always exists a totally odd part-square decomposition $h = q^2 + \rho$ with $q(z), \rho(z) \in \bar{K}[z]$.

Proof. We write $h_e(z)$ and $h_o(z)$ for the sums of the even- and odd- degree terms of $h(z)$ respectively, so that $h = h_e + h_o$. We denote the degree of h_e by $2m$. As all terms of h_e have even degree, we may write $h_e(z) = \hat{h}(z^2)$ for some uniquely determined $\hat{h}(z) \in K[z]$ of degree m. Let $\alpha_1, \ldots, \alpha_m \in \bar{K}$ be the roots of $\hat{h}$, and let us denote by $c \in K$ its leading coefficient. Let us also choose a square root $\sqrt{c} \in \bar{K}$ of c and a square root $\sqrt{\alpha_i} \in \bar{K}$ of each root of $\hat{h}$, and let us define

$$\begin{aligned}\hat{h}_+(z) :=& \sqrt{c}\prod_i (z + \sqrt{\alpha_i}) = c_0 z^m + c_1 z^{m-1} + \ldots + c_m,\\ \hat{h}_-(z) :=& \sqrt{c}\prod_i (z - \sqrt{\alpha_i}) = c_0 z^m - c_1 z^{m-1} + \ldots + (-1)^m c_m.\end{aligned}$$

It is clear that we have $h_e(z) = \hat{h}_+(z)\hat{h}_-(z)$; exploiting this factorization of h_e, we may write the kth-order coefficient of h, whenever k is even, as

$$\sum_{i+j=2m-k} (-1)^i c_i c_j. \tag{15}$$

If we now choose a square root $\sqrt{-1} \in \bar{K}$ of -1, and we define

$$c_i' := \begin{cases} c_i & \text{if } i \text{ is even,} \\ \sqrt{-1} \cdot c_i & \text{if } i \text{ is odd,} \end{cases}$$

we may rewrite the expression (15), for all even values of k, in the more symmetric form

$$\sum_{i+j=2m-k} c_i' c_j'. \tag{16}$$

If we now set $q(z) = c_0' z^m + \ldots + c_m'$, the even-degree terms of q^2 reproduce h_e. Hence, the polynomial $\rho(z) := h(z) - q^2(z)$ only consists of odd-degree terms: in other words, $h = q^2 + \rho$ is a totally odd part-square decomposition for h. $\square$

We note that, if a nonzero polynomial $h \in \bar{K}[z]$ is written as a product of factors $h = \prod_{i=1}^N h_i$ with $h_i \in \bar{K}[z]$, then, given part-square decompositions $h_i = q_i^2 + \rho_i$, with $q_i, \rho_i \in \bar{K}[z]$, one can use them to form a part-square decomposition $h = q^2 + \rho$, where $q = \prod_{i=1}^N q_i$ and $\rho = h - q^2$. We have the following.

Proposition 4.21. In the setting above, let $t_i := t_{q_i,\rho_i}$ and $t := t_{q,\rho}$.

(a) If $t_i \geq 0$ for all i, then we have $t \geq \min\{t_1, \ldots t_N\}$.

(b) If $t_i \geq 0$ for all i, and the minimum $\min\{t_1, \ldots t_N\}$ is achieved by only one of the t_i's, then we have $t = \min\{t_1, \ldots t_N\}$; moreover, in this case, if i_0 is the index such that $t_{i_0} < t_i$ for all i, then the part-square decomposition of h is good if and only if that of h_{i_0} is.

(c) Assume that $N = 2$, and suppose that, for all roots s_1 in $\bar{K}$ of h_1 and for all roots s_2 of h_2, we have $v(s_1) > 0$ but $v(s_2) < 0$; assume, moreover, that both decompositions $h_i = q_i^2 + \rho_i$ are good. Then, if $\min\{t_1, t_2\} < 2v(2)$, we have $t = \min\{t_1, t_2\}$, and the corresponding decomposition of h is also good.

Proof. Let us first address points (a) and (b). It is clearly enough to prove these results for $N = 2$. In this case, we have

$$\rho = h - q^2 = (q_1^2 + \rho_1)(q_2^2 + \rho_2) - (q_1 q_2)^2 = g_1 + g_2 + g_3, \tag{17}$$

where $g_1, g_2, g_3 \in \bar{K}[z]$ are the polynomials

$$g_1 = \rho_1 q_2^2, \qquad g_2 = \rho_2 q_1^2, \qquad g_3 = \rho_1 \rho_2. \tag{18}$$

Since $t_i \geq 0$, i.e. $v(\rho_i) \geq v(h_i)$, we have $v(q_i^2) = v(h_i - \rho_i) \geq v(h_i)$ for $i = 1, 2$. Moreover, we have $v(\rho_i) = v(h_i) + t_i$. We therefore get $v(g_1) \geq t_1 + v(h), v(g_2) \geq t_2 + v(h)$, and $v(g_3) \geq t_1 + t_2 + v(h)$; all three thresholds are clearly $\geq \min\{t_1, t_2\} + v(h)$, from which we deduce $v(\rho) \geq \min\{t_1, t_2\} + v(h)$, and thus $t \geq \min\{t_1, t_2\}$.

To prove part (b), let us now further assume that $t_1 < t_2$. This implies, in particular, $t_2 > 0$; hence we have $v(\rho_2) > v(h_2)$ and $v(q_2^2) = v(h_2 - \rho_2) = v(h_2)$, and we consequently get $v(g_1) = t_1 + v(h)$. We deduce from this that $v(g_2) \geq t_2 + v(h) > v(g_1)$, and $v(g_3) \geq t_1 + t_2 + v(h) > v(g_1)$. It follows that $v(\rho) = v(g_1) = t_1 + h$, implying that $t = t_1 = \min\{t_1, t_2\}$. Moreover, the normalized reduction of ρ equals that of $g_1 = \rho_1 q_2^2$; as a consequence, the normalized reduction of ρ is a square if and only if that of ρ_1 is. From this, together with the fact that $t_1 = t$, we deduce that $h = q^2 + \rho$ is a good decomposition if and only if $h_1 = q_1^2 + \rho_1$ is (see Proposition 4.18).

Let us now address part (c). Let d_i be the degree of h_i, and let $\gamma_i \in \bar{K}^\times$ be an element of valuation $v(h_i)$ for $i = 1, 2$. Since we are assuming that $v(s_1) > 0$, for all root s_1 of h_1 we have $\overline{\gamma_1^{-1} h_1}(z) = c_1 z^{d_1}$ for some $c_1 \in k^\times$; similarly, since we have $v(s_2) < 0$ for all roots s_2 of h_2 we have that $\overline{\gamma_2^{-1} h_2}(z)$ is a constant $c_2 \in k^\times$. Since the decompositions of h_1 and h_2 are assumed to be good, we have $t_1, t_2 \geq 0$; moreover, from the fact that the normalized reduction of h_2 is a square we deduce, via Proposition 4.18, that $t_2 > 0$. Now, when $t_1 \neq t_2$ the conclusion follows from (b). We are consequently only left to address the case where $0 < t_1 = t_2 < 2v(2)$. Here we already know that $t > 0$ from part (a), and that d_1 must necessarily be even, because, since $t_1 > 0$, $\overline{\gamma_1^{-1} h_1} = \overline{\gamma_1^{-1} q_1^2}$ must be a square; let us write $d_1 = 2m$.

In the case we are considering, we clearly have that $v(g_1) = v(g_2) = t_1 + v(h)$, while $v(g_3) > t_1 + v(h)$. If we let $\gamma \in \bar{K}^\times$ be any element of valuation t, we consequently have that

$$\overline{\gamma_1^{-1} \gamma_2^{-1} \gamma^{-1} \rho} = (\overline{\gamma_2^{-1} h_2}) \overline{r_1} + (\overline{\gamma_1^{-1} h_1}) \overline{r_2} = c_2 \overline{r_1} + c_1 z^{2m} \overline{r_2} \tag{19}$$

where $r_1 := \gamma^{-1}\gamma_1^{-1}\rho_1$ and $r_2 := \gamma^{-1}\gamma_2^{-1}\rho_2$, so that $\overline{r}_1$ and $\overline{r}_2$ are normalized reduction of ρ_1 and ρ_2, respectively. We remark $\overline{r_1}$ has degree $\deg(\overline{r}_1) \leq 2m$; hence, if an odd-degree term of degree s appears in the normalized reduction $\overline{r_1}$ of ρ_1 (resp. in the normalized reduction $\overline{r_2}$ of ρ_2), then an odd-degree term of degree s (resp. $s + 2m$) will also show up in $\overline{\gamma_1^{-1}\gamma_2^{-1}\gamma^{-1}\rho}$: roughly speaking, in the expression for $\overline{\gamma_1^{-1}\gamma_2^{-1}\gamma^{-1}\rho}$ no cancellation occurs between the odd-degree monomials of $\overline{r_1}$ and those of $\overline{r_2}$. We conclude that, since $\overline{r_1}$ and $\overline{r_2}$ are not squares by Proposition 4.18, the reduced polynomial $\overline{\gamma_1^{-1}\gamma_2^{-1}\gamma^{-1}\rho}$ is not a square, so that $v(\rho) = t_1 + v(h)$ (i.e., $t = t_1 = t_2$), and the decomposition of h is good by Proposition 4.18. $\square$

4. Forming models of Y using part-square decompositions

In this subsection, we compute the model of the hyperelliptic curve $Y : y^2 = f(x)$ corresponding to any given smooth model of the projective line X, in the sense of §3.1. More precisely, let $D := D_{\alpha,b}$ be a disc in $\bar{K}$ with $\alpha \in \bar{K}$ and $b = v(\beta)$ for some $\beta \in \bar{K}^\times$. To this disc we can attach (see §4.2) a smooth model $\mathcal{X}_D$ of the line X, defined over some extension of R. We will show that, after possibly replacing this extension with a further extension R', which in particular will be large enough so that $f_{\alpha,\beta}$ admits a good part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ over the fraction field K' of $\mathrm{Frac}(R')$, the model of Y corresponding to $\mathcal{X}_D/R'$ has reduced special fiber, and its equation can explicitly be written using $q_{\alpha,\beta}$ and $\rho_{\alpha,\beta}$; we will denote this model by $\mathcal{Y}_D$. Here, $x_{\alpha,\beta}$ denotes the coordinate obtained from x by translating by α and scaling by β, and $f_{\alpha,\beta}(x_{\alpha,\beta})$ is the polynomial obtained by applying those transformations to f (see §1.5.5.4 for this notational convention).

The strategy will be the following one: after a suitable change of the coordinate y, we will rewrite the equation $y^2 = f_{\alpha,\beta}(x_{\alpha,\beta})$ of the hyperelliptic curve Y in the form

$$y^2 + q_0(x_{\alpha,\beta})y - \rho_0(x_{\alpha,\beta}) = 0, \quad \text{with } \deg(q_0) \leq g+1, \ \deg(\rho_0) \leq 2g+2, \tag{20}$$

such that the following conditions are satisfied:

(a) ρ_0 and q_0 have integral coefficients (i.e., we have $\rho_0(x_{\alpha,\beta}), q_0(x_{\alpha,\beta}) \in R'[x_{\alpha,\beta}]$);
(b) the k-curve given by the reduction of the equation in (20) is reduced.

Then, the model $\mathcal{Y}_D$ is constructed as follows. The equation in (20) above defines a scheme W over R' whose generic fiber is isomorphic to the affine chart $x_{\alpha,\beta} \neq \infty$ of the hyperelliptic curve Y. The coordinate $x_{\alpha,\beta}$ defines a map $W \to \mathcal{X}_D$, whose image is the affine chart $x_{\alpha,\beta} \neq \infty$ of $\mathcal{X}_D$. Over the affine chart $x_{\alpha,\beta} \neq 0$ of $\mathcal{X}_D$, we can correspondigly form the R'-scheme $W^\vee$ defined by the equation

$$\begin{aligned} &\check{y}^2 + q_0^\vee(\check{x}_{\alpha,\beta})\check{y} - \rho_0^\vee(\check{x}_{\alpha,\beta}) = 0, \qquad \text{where } \check{x}_{\alpha,\beta} = x_{\alpha,\beta}^{-1}, \ \check{y} = x_{\alpha,\beta}^{-(g+1)}y, \\ &\qquad q_0^\vee(\check{x}_{\alpha,\beta}) = {x_{\alpha,\beta}}^{-(g+1)}q_0(x_{\alpha,\beta}), \text{ and } \rho_0^\vee(\check{x}_{\alpha,\beta}) = {x_{\alpha,\beta}}^{-(2g+2)}\rho_0(x_{\alpha,\beta}). \end{aligned} \tag{21}$$

We can now *define* $\mathcal{Y}_D$ to be the scheme obtained by gluing the affine charts W and $W^\vee$ together in the obvious way: it is endowed with a degree-2 covering map $\mathcal{Y}_D \to \mathcal{X}_D$, and its generic fiber is identified with the hyperelliptic curve $Y \to X$.

Proposition 4.22. The scheme $\mathcal{Y}_D$ constructed above, which is defined over an appropriate extension R' of R, coincides with the normalization of $\mathcal{X}_D/R'$ in the function field of the hyperelliptic curve Y, and it is a model of Y whose special fiber is reduced.

PROOF. We have to show that the scheme $\mathcal{Y}_D$ we have constructed is normal. The R'-schemes W and $W^\vee$ are complete intersections, and hence they are Cohen-Macaulay; as a consequence, to check that $\mathcal{Y}_D$ is normal, it is enough to prove that it is regular at its codimension-1 points. Since the generic fiber of $\mathcal{Y}_D$ coincides with Y, it is certainly regular; hence, all that is left is to check that $\mathcal{Y}_D$ is regular at the generic point η_{V_i} of each irreducible component V_i of the special fiber $(\mathcal{Y}_D)_s$. Since we are assuming that the k-curve W_s is reduced, the k-curve $(\mathcal{Y}_D)_s$ is also clearly reduced, which implies that $\mathcal{Y}_D$ is certainly regular at the points η_{V_i}. Thus, the scheme $\mathcal{Y}_D$ is actually normal.

It is now completely clear that $\mathcal{Y}_D$ is the model of Y obtained by normalizing $\mathcal{X}_D/R'$ in the function field of Y. □

All we have to do now is determine a change of the coordinate y such that conditions (a) and (b) above are satisfied. To do this, suppose that we are given a good part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ (which certainly exists over some extension of K, thanks to Proposition 4.20 and Corollary 4.19). Let $\gamma \in \bar{K}^\times$ be an element whose valuation is $v(\gamma) = \min\{t, 2v(2)\} + v(f_{\alpha,\beta})$, where $t = t_{q_{\alpha,\beta},\rho_{\alpha,\beta}} = v(\rho_{\alpha,\beta}) - v(f_{\alpha,\beta})$. We remark that we necessarily have $t \geq 0$, since the part-square decomposition is assumed to be good (see Remark 4.16). The change of variable we perform is $y \mapsto \gamma^{1/2}y + q_{\alpha,\beta}(x_{\alpha,\beta})$, and it leads to an equation of the form (20) with

$$q_0 = 2\gamma^{-1/2}q_{\alpha,\beta} \qquad \text{and} \qquad \rho_0 = \gamma^{-1}\rho_{\alpha,\beta}. \tag{22}$$

The valuations of q_0 and ρ_0 can be computed as follows.

(1) For q_0, we have $2v(q_0) = 2v(2) - \min\{t, 2v(2)\} + 2v(q_{\alpha,\beta}) - v(f_{\alpha,\beta})$. Let us remark that, since $t \geq 0$, we have $2v(q_{\alpha,\beta}) \geq v(f_{\alpha,\beta})$, and moreover equality holds whenever $t > 0$. We deduce that:
   (i) $v(q_0) \geq 2v(2) - \min\{t, 2v(2)\}$ for all t, so that q_0 is consequently always integral;
   (ii) $v(q_0) = 2v(2) - \min\{t, 2v(2)\}$ whenever $t > 0$;
   (iii) $v(q_0) > 0$ if $0 \leq t < 2v(2)$ (which can only happen in the $p = 2$ setting); and
   (iv) $v(q_0) = 0$ if $t \geq 2v(2)$ in the $p = 2$ setting.
(2) For ρ_0, we have $v(\rho_0) = t - \min\{t, 2v(2)\}$; in particular,
   (i) ρ_0 is always integral;
   (ii) $v(\rho_0) = 0$ if $0 \leq t \leq 2v(2)$; and
   (iii) $v(\rho_0) > 0$ if $t > 2v(2)$.

These computations guarantee that condition (a) is satisfied. We now verify that also condition (b) is satisfied.

Lemma 4.23. In the context above, condition (b) is also satisfied, i.e. the reduction of equation (20) defines a reduced k-curve. Moreover, this curve is a separable (resp. inseparable) quadratic cover of the k-line of coordinate $x_{\alpha,\beta}$ if and only if $t \geq 2v(2)$ (resp. $0 \leq t < 2v(2)$).

PROOF. Suppose by way of contradiction that the k-curve defined by the reduction of (20) is non-reduced. This is clearly equivalent to saying that the polynomial $g(x_{\alpha,\beta}, y) \in k[x_{\alpha,\beta}, y]$ given by the reduction of (20) (i.e. $g(x_{\alpha,\beta}, y) := y^2 + \overline{q_0(x_{\alpha,\beta})}y - \overline{\rho_0(x_{\alpha,\beta})}$) is a square. If we treat $g(x_{\alpha,\beta}, y)$ as a monic quadratic polynomial in the variable y, we can say that it is a square if and only if its constant term $\overline{\rho_0(x_{\alpha,\beta})} \in k[x_{\alpha,\beta}]$ is a square, and its discriminant $\Delta = \overline{q_0}^2 + 4\overline{\rho_0} = \overline{4\gamma^{-1} f_{\alpha,\beta}} \in k[x_{\alpha,\beta}]$ is zero. However, when $t \geq 2v(2)$, we have $v(\gamma) = v(4f_{\alpha,\beta})$ and therefore $\Delta \neq 0$; when $0 \leq t < 2v(2)$, the reduced polynomial $\overline{\rho_0}$ is a normalized reduction of $\rho_{\alpha,\beta}$, which is not a square by Proposition 4.18. We conclude that the k-curve $g(x_{\alpha,\beta}, y) = 0$ is always reduced.

Now the coordinate $x_{\alpha,\beta}$ defines a quadratic cover from the k-curve $g(x_{\alpha,\beta}, y) = 0$ to the affine k-line, and it is immediate to realize that this cover is inseparable only when $p = 2$ and the linear term $\overline{q_0(x_{\alpha,\beta})}y$ vanishes, which happens if and only if $0 < t \leq 2v(2)$. □

The following proposition summarizes the results we have obtained.

Proposition 4.24. Let $\mathcal{X}_D$ be the smooth model of the line corresponding to the disc $D := D_{\alpha,v(\beta)}$, with $\alpha \in \bar{K}$ and $\beta \in \bar{K}^\times$. Then, after replacing K with an appropriate finite extension, the normalization $\mathcal{Y}_D$ of $\mathcal{X}_D$ in $K(Y)$ has reduced special fiber. Given a good part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$, and letting $t = t_{q_{\alpha,\beta},\rho_{\alpha,\beta}}$, the model $\mathcal{Y}_D$ falls under (exactly) one of the following two cases:

(1) $t \geq 2v(2)$; in this case, $(\mathcal{Y}_D)_s$ is a separable degree-2 cover of $(\mathcal{X}_D)_s$; and
(2) $0 \leq t < 2v(2)$; in this case, $(\mathcal{Y}_D)_s$ is an inseparable degree-2 cover of $(\mathcal{X}_D)_s$.

The equations describing the affine charts $x_{\alpha,\beta} \neq \infty$ and $x_{\alpha,\beta} \neq 0$ of the model $\mathcal{Y}_D$ have the form (20) and (21) respectively, and they can be explicitly computed from $q_{\alpha,\beta}$ and $\rho_{\alpha,\beta}$ using the formulas in (22).

5. The special fiber $(\mathcal{Y}_D)_s$ in the separable case

We will now study the special fiber of the model $\mathcal{Y}_D$ associated to a given disc $D := D_{\alpha,v(\beta)}$ be a disc with $a \in \bar{K}$ and $\beta \in \bar{K}^\times$, which was computed in the previous subsection. This subsection will consider the case in which $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is separable: this means that it is possible to find a part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ satisfying $t := t_{q_{\alpha,\beta},\rho_{\alpha,\beta}} \geq 2v(2)$, and the equation of $(\mathcal{Y}_D)_s$ has the form $y^2 + \overline{q_0(x_{\alpha,\beta})}y = \overline{\rho_0(x_{\alpha,\beta})}$, where $q_0 := 2\gamma^{-1/2}q_{\alpha,\beta}$, and $\rho_0 = \gamma^{-1}\rho_{\alpha,\beta}$, where $\gamma \in \bar{K}^\times$ is an element of valuation $v(f_{\alpha,\beta}) + 2v(2)$. In the $p \neq 2$ case, the equation of $(\mathcal{Y}_D)_s$ can also be written in the simpler form $y^2 = \overline{f_0(x_{\alpha,\beta})}$, where

$$f_0 = 4\gamma^{-1} f_{\alpha,\beta} = q_0^2 + 4\rho_0. \tag{23}$$

We remark that the separable quadratic cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is branched precisely above the points $P_1, \ldots, P_N$ of $(\mathcal{X}_D)_s$ at which the roots $\mathcal{R} \cup \{\infty\}$ reduce and is étale elsewhere: this can be seen directly from the equation of $(\mathcal{Y}_D)_s$, or can be deduced from the fact the branch locus of $\mathcal{Y}_D \to \mathcal{X}_D$ has pure dimension 1 by Zariski–Nagata purity theorem. In order to state and prove the results in this subsection, we partition the branch locus $R = \{P_1, \ldots, P_N\} \subseteq (\mathcal{X}_D)_s(k)$ in three subsets as $R = R_0 \sqcup R_1 \sqcup R_2$, in the following way.

$R_0 = \{P \in (\mathcal{X}_D)_s : (\mathcal{Y}_D)_s$ exhibits a unique smooth point Q above $P\}$;

$R_1 = \{P \in (\mathcal{X}_D)_s : (\mathcal{Y}_D)_s$ has a (unique) singular point Q above P and has one branch at $Q\}$;

$R_2 = \{P \in (\mathcal{X}_D)_s : (\mathcal{Y}_D)_s$ has a (unique) singular point Q above P and has two branches at $Q\}$.

We denote the cardinality of each subset $R_i \subseteq R$ by N_i for $i = 0, 1, 2$.

Remark 4.25. The following statements are clear from the definitions above.

(a) The set $R_0 \cup R_1$ is precisely the branch locus of the quadratic cover $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$, where $\widetilde{(\mathcal{Y}_D)_s}$ is the normalization of the k-curve $(\mathcal{Y}_D)_s$.

(b) The curve $(\mathcal{Y}_D)_s$ has exactly $N_1 + N_2$ singular points, which lie over the $N_1 + N_2$ points of $R_1 \cup R_2$.

(c) The unique point $Q \in (\mathcal{Y}_D)_s$ lying over some given $P \in R$ is fixed by the action of the hyperelliptic involution. If $P \in R_2$, the two branches of $(\mathcal{Y}_D)_s$ passing through Q get flipped by the hyperelliptic involution.

(d) The special fiber $(\mathcal{Y}_D)_s$ either consists of two components flipped by the hyperelliptic involution, or it is irreducible. In the first case (which always occurs, for example, if $\overline{\rho_0(x_{\alpha,\beta})}$ is the zero polynomial, i.e. if $t > 2v(2)$), the two components are necessarily two lines that trivially cover $(\mathcal{X}_D)_s$, while $\widetilde{(\mathcal{Y}_D)_s}$ is their disjoint union, and we have $R_0 \cup R_1 = \varnothing$. If $(\mathcal{Y}_D)_s$ is irreducible, however, the quadratic cover $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ is necessarily ramified, because $\mathbb{P}^1_k$ does not have non-trivial finite étale connected covers: hence, and we have $R_0 \cup R_1 \neq \varnothing$.

We want to better understand the ramification behaviour of $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ above the points of $R_0 \cup R_1$; to this aim, we can measure, above each point, the length of the module of relative Kähler differentials of the cover.

Definition 4.26. Given $P \in (\mathcal{X}_D)_s(k)$, we set

$$\ell(\mathcal{X}_D, P) = \mathrm{length}_{\mathcal{O}_{(\mathcal{X}_D)_s,P}} \left(\Omega_{\widetilde{(\mathcal{Y}_D)_s}/(\mathcal{X}_D)_s} \otimes \mathcal{O}_{(\mathcal{X}_D)_s,P} \right).$$

Remark 4.27. For any $P \in (\mathcal{X}_D)_s(k)$, the integer $\ell(\mathcal{X}_D, P)$ satisfies the following properties.

(a) If $P \notin R_0 \cup R_1$, then $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ is unramified over P, and we thus have $\ell(\mathcal{X}_D, P) = 0$.

(b) If $P \in R_0 \cup R_1$, and we denote by Q its unique preimage $\widetilde{(\mathcal{Y}_D)_s}$, the ramification index of the cover $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ at Q is $e_Q = 2$, and [11, Proposition 7.4.13] ensures that $\ell(\mathcal{X}_D, P) \geq e_Q - 1$, with equality if and only if the cover is tame. This means that $\ell(\mathcal{X}_D, P) = 1$ if $p \neq 2$, and $\ell(\mathcal{X}_D, P) \geq 2$ if $p = 2$.

The knowledge of ℓ at the points of $(\mathcal{X}_D)_s$ gives us information about the abelian rank of $(\mathcal{Y}_D)_s$.

Proposition 4.28. The genus of $\widetilde{(\mathcal{Y}_D)_s}$ is given by

$$g\left(\widetilde{(\mathcal{Y}_D)_s}\right) = -1 + \frac{1}{2} \sum_{P \in (\mathcal{X}_D)_s(k)} \ell(\mathcal{X}_D, P), \tag{24}$$

with the convention that the genus of the disjoint union of two lines is -1. In the $p \neq 2$ setting, this can be rewritten as

$$g\left(\widetilde{(\mathcal{Y}_D)_s}\right) = -1 + \frac{1}{2}(N_0 + N_1), \tag{25}$$

while, in the $p = 2$ setting, the formula in (24) implies the inequality

$$g\left(\widetilde{(\mathcal{Y}_D)_s}\right) \geq -1 + (N_0 + N_1). \tag{26}$$

PROOF. Equation (24) is just the Riemann-Hurwitz formula (see, for example, [11, Theorem 7.4.16]), while (25) and (26) follow from (24) via Remark 4.27. □

Remark 4.29. In particular, the formula in (25) implies that, in the $p \neq 2$ setting, the integer $N_0 + N_1$ is necessarily even.

We now see how to compute $\ell(\mathcal{X}_D, P)$ for a given point $P \in (\mathcal{X}_D)_s$ from the good part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ given.

Lemma 4.30. Choose $P \in (\mathcal{X}_D)_s(k)$. Let us denote by $n_q(P) := \mathrm{ord}_P(\overline{q_0})$, $n_\rho(P) := \mathrm{ord}_P(\overline{\rho_0})$, $n_f(P) := \mathrm{ord}_P(\overline{f_0})$ the respective orders of vanishing at the point P of the reductions of the polynomials q_0, ρ_0 and f_0 defined in (22) and (23), with the convention that the zero polynomial has vanishing order ∞, and that, if $P = \infty$, the vanishing orders of $\overline{f_0}$, $\overline{\rho_0}$ and $\overline{q_0}$ at ∞ are respectively those of $\overline{f_0^\vee}$, $\overline{\rho_0^\vee}$ and $\overline{q_0^\vee}$ at 0, i.e. $n_q = g+1-\deg(\overline{q_0})$, $n_\rho = 2g+2-\deg(\overline{\rho_0})$ and $n_f = 2g+2-\deg(\overline{f_0})$. Then,

(a) if $p \neq 2$, then $\ell(\mathcal{X}_D, P)$ is 0 (resp. 1) if the integer n_f is even (resp. odd);
(b) if $p = 2$ and $2n_q(P) \leq n_\rho(P)$, then we have $\ell(\mathcal{X}_D, P) = 0$; and
(c) if $p = 2$ and if $2n_q(P) > n_\rho(P)$ and $n_\rho(P)$ is odd, then we have $\ell(\mathcal{X}_D, P) = 2n_q(P) - n_\rho(P) + 1$.

PROOF. We lose no generality in assuming P has coordinate $\overline{x_{\alpha,\beta}} = 0$. For brevity we write z for the variable $x_{\alpha,\beta}$. We proceed by desingularizing $(\mathcal{Y}_D)_s$ above P by means of a sequence of blowups. Let us first work in the $p \neq 2$ setting. The equation of $(\mathcal{Y}_D)_s$, in this case, has the form $y^2 = \overline{f_0}(z) = z^{n_f} f_1(z)$, with $f_1(z) \in k[z]$ and $f_1(0) \neq 0$. If $n_f = 0$, then we are already done; otherwise, the curve becomes nonsignular above $z = 0$ after blowing it up $\lfloor n_f/2 \rfloor$ times at $(0,0)$; at each blowup, the right-hand side of the equation is divided by z^2, so that the desingularized equation becomes $y^2 = z^e f_1(z)$, where e is 0 or 1, depending on whether n_f is even or odd; moreover, when $e = 1$ this curve is ramified over $z = 0$, whereas, when $e = 0$, it is étale over $z = 0$. From this, (a) follows, taking into account Remark 4.27.

Let us now adopt the $p = 2$ setting. The equation of $(\mathcal{Y}_D)_s$ is now $y^2 + \overline{q_0}(z)y = \overline{\rho_0}(z)$, with $\overline{q_0}(z) = z^{n_q} q_1(z)$, and $\overline{\rho_0}(z) = z^{n_\rho} r_1(z)$, where $q_1(z), \rho_1(z) \in k[z]$ do not vanish at 0.

Assume that $2n_q \leq n_\rho$. Then, after n_q blowups at $(0,0)$, we obtain $y^2 + q_1(z)y = z^{n_\rho - 2n_q}\rho_1(z)$. Since $q_1(0) \neq 0$, there are exactly 2 solutions for y at $z = 0$, which means that the blown-up curve is étale above P, implying that $\ell(\mathcal{X}_D, P) = 0$. We have thus proved part (b).

Assume that $2n_q > n_\rho$ and that n_ρ is odd. Then, after $(n_\rho - 1)/2$ blowups at $(0,0)$, we obtain the equation

$$y^2 + z^{n_q - (n_\rho - 1)/2} q_1(z) y = z\rho_1(z). \tag{27}$$

The curve given by (27) has a unique point $(0,0)$ above $z = 0$ and it is non-singular at that point; this is enough to guarantee that $\ell(\mathcal{X}_D, P) > 0$. Let $B := k[z,y]_{(z)}/(\text{equation in (27)})$ be the local ring of functions on the blown-up curve at $(0,0)$, which is a free $k[z]_{(z)}$-algebra of rank 2. Then $\ell(\mathcal{X}_D, P)$ equals the length of the $k[z]_{(z)}$-module $\Omega_{B/k[z]_{(z)}}$, or, equivalently, the dimension over k of $\Omega_{B/k[z]_{(z)}}$. We have an isomorphism of $k[z]_{(z)}$-modules

$$\Omega_{B/k[z]_{(z)}} = Bdy/(z^{n_q-(n_\rho-1)/2}q_1(z)dy) \xrightarrow{\sim} (k[z]_{(z)})[y]/(y^2 - zr_1(z), z^{n_q-(n_\rho-1)/2}q_1(z)), \tag{28}$$

where the isomorphism is given by sending dy to 1. The latter $k[z]_{(z)}$-module, however, is a free algebra of rank 2 over the ring

$$k[z]_{(z)}/(z^{n_q-(n_\rho-1)/2}q_1(z)) \cong k[z]/(z^{n_q-(n_\rho-1)/2}),$$

which clearly has dimension $n_q - (n_\rho - 1)/2$ over k. From this, part (c) follows. □

Remark 4.31. We make the following observations about the subsets $R_i \subseteq R$.

(a) Assume that $p \neq 2$. Lemma 4.30 tells us that, for all $P \in (\mathcal{X}_D)_s$, the integer $\ell(\mathcal{X}_D, P)$ is 0 or 1 depending on whether an even or an odd number of the $2g+2$ points of $\mathcal{R} \cup \{\infty\}$ reduce to P. In light of Remark 4.27, we conclude that, in the $p \neq 2$ case, R_2 (resp. $R_0 \cup R_1$) is the set of points of $(\mathcal{X}_D)_s$ at which an *even* (resp. *odd*) number of roots of $\mathcal{R} \cup \{\infty\}$ reduce. Actually, it is also easy to see that
   (i) $P \in R_0$ if and only if exactly one root of $\mathcal{R} \cup \{\infty\}$ reduces to it;
   (ii) $P \in R_1$ if and only if only an odd number ≥ 3 of roots of $\mathcal{R} \cup \{\infty\}$ reduces to it; and
   (iii) $P \in R_2$ if and only if an even number ≥ 2 of roots of $\mathcal{R} \cup \{\infty\}$ reduces to it.

   When we partition the even-cardinality set $\mathcal{R} \cup \{\infty\}$ according to the points of $(\mathcal{X}_D)_s$ at which its elements reduce, the number of odd cardinality classes must be even: this shows that $N_1 + N_0$ is even, as we have already observed in Remark 4.29.

(b) Assume that $p = 2$. Lemma 4.30 allows us to calculate $\ell(\mathcal{X}_D, P)$ from a given good part-square decomposition of $f_{\alpha,\beta}$ only in certain cases: in fact, when $2n_q(P) > n_\rho(P)$ and $n_\rho(P)$ is even, the lemma is inconclusive. At the same time, we remark that if we choose a *totally odd* part-square decomposition for $f_{\alpha,\beta}$ (which can always be done by Proposition 4.20), the polynomial $\overline{\rho_0}$ will certainly have a zero of odd multiplicity at the points 0 and ∞ of $(\mathcal{X}_D)_s$; hence, we will certainly be able to compute ℓ at the points $\overline{x_{\alpha,\beta}} = 0$ and $\overline{x_{\alpha,\beta}} = \infty$ via the lemma. In other words, given a point $P \in (\mathcal{X}_D)_s$, by appropriately choosing the center α of the disc D and constructing a totally odd decomposition for $f_{\alpha,\beta}$, Lemma 4.30 allows us to compute $\ell(\mathcal{X}_D, P)$ at the point, and the result it produces is a non-negative even integer.

We now give a criterion to determine whether $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ (which is equivalent to saying that $\mathcal{Y}_D \leq \mathcal{Y}^{\mathrm{rst}}$).

Theorem 4.32. Assume that D is a disc such that $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is separable, and let N denote the number of points of $(\mathcal{X}_D)_s$ to which the roots $\mathcal{R} \cup \{\infty\}$ reduce. We have $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ if and only if one of the following conditions holds:

(1) $N \geq 3$;
(2) $p = 2$, $N = 2$ and $(\mathcal{Y}_D)_s$ is irreducible; or
(3) $p = 2$, $N = 1$ and $(\mathcal{Y}_D)_s$ is irreducible of positive abelian rank.

Moreover, whenever $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$, the strict transform of the k-curve $(\mathcal{Y}_D)_s$ in $(\mathcal{Y}^{\mathrm{rst}})_s$ is smooth, and it consequently coincides with its normalization $\widetilde{(\mathcal{Y}_D)_s}$.

Proof. The result essentially follows from a combinatorial argument that directly makes use of the description we have given of $(\mathcal{Y}_D)_s$ in this subsection, by applying the criterion we have presented in Proposition 3.17. Let us write $N = N_0 + N_1 + N_2$ as we did at the beginning of this subsection; we recall that the integers N_0, N_1, and N_2 are respectively the number of points of $(\mathcal{X}_D)_s$ above which $(\mathcal{Y}_D)_s$ is ramified and exhibits a smooth point, a singular point through which only one branch of $(\mathcal{Y}_D)_s$ passes, and a singular point through which two branches of $(\mathcal{Y}_D)_s$ pass. We also recall from Remark 4.29 that, in the $p \neq 2$ setting, the integer $N_0 + N_1$ is necessarily even.

Suppose that $(\mathcal{Y}_D)_s$ is not irreducible. As we have seen in Remark 4.25(d), this is equivalent to the saying that $N_0 = N_1 = 0$, and the curve $(\mathcal{Y}_D)_s$ consists, in this case, of two lines L_1 and L_2 meeting each other above the $N = N_2$ points of $(\mathcal{X}_D)_s$; the number of singular points of $(\mathcal{Y}_D)_s$ is N, and through each singular point one branch of L_1 and one branch of L_2 pass, flipped by the hyperlliptic involution. We have $m(L_i) = 1$, $a(L_i) = 0$, $w(L_i) = N$, and $\underline{w}(L_i) = (1, \ldots, 1)$ for $i = 1, 2$; hence, Proposition 3.17 ensures that $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ if and only if $N \geq 3$.

Suppose now that $(\mathcal{Y}_D)_s$ is irreducible, which is to say that $N_0 + N_1 \geq 1$, and let $V = (\mathcal{Y}_D)_s$ denote the unique irreducible component of $(\mathcal{Y}_D)_s$. We have the following:

- $w(V) = N_1 + 2N_2$;
- $\underline{w}(V) = (1, \ldots, 1, 2, \ldots, 2)$ with 1 appearing N_1 times and 2 appearing N_2 times; and
- $a(V) = -1 + (N_0 + N_1)/2$ in the $p \neq 2$ case, and $a(V) \geq -1 + N_0 + N_1$ in the $p = 2$ setting, by Proposition 4.28.

Suppose that $N = 1$. Then we have $N_0 + N_1 = 1$, which is impossible if $p \neq 2$ (as it contradicts Remark 4.29) and so we must have $p = 2$. If $a(V) \geq 1$ then by Proposition 3.17 we have $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$, while if $a(V) = 0$, then we have $w(V) \leq 1$ and so Proposition 3.17 says that $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$.

Suppose now that $N = 2$ and $p \neq 2$. This forces $N_0 + N_1 = 2$ by Remark 4.29, from which it follows that $a(V) = 0$; meanwhile, we have $w(V) = N_1 + 2N_2 = N_1 \leq 2$ and $\underline{w}(V)$ consists only of 1's, and so by Proposition 3.17 we have $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$.

Finally, suppose that $N \geq 3$ or that $N = 2$ and $p = 2$. If $N_2 \geq 1$, then we have $w(V) \geq 2$ and that a 2 appears in $\underline{w}(V)$, and so $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ by Proposition 3.17. If $N_2 = 0$, then we have $N_0 + N_1 \geq 4$ if $p \neq 2$ by Remark 4.29 and $N_0 + N_1 \leq 2$ if $p = 2$; either way, we get $a(V) \geq 1$, and so again $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ by Proposition 3.17.

The statement about the strict transform of $(\mathcal{Y}_D)_s$ in $(\mathcal{Y}^{\mathrm{rst}})_s$ is an immediate consequence of Proposition 3.16, taking into account that the irreducible components of the special fiber of a semistable model of the line are always lines, and hence, in particular, smooth k-curves. $\square$

6. The special fiber $(\mathcal{Y}_D)_s$ in the inseparable case

We again let $D := D_{\alpha, v(\beta)}$ be a disc with $a \in \bar{K}$ and $\beta \in \bar{K}^{\times}$, and let $\mathcal{Y}_D$ be the corresponding model of Y constructed in §4.4; this subsection will analyze the case in which $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable (we are thus in the $p = 2$ setting). In this case, given a good part-square decomposition $f_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$, we have $0 \leq t := t_{q_{\alpha,\beta},\rho_{\alpha,\beta}} < 2v(2)$, and the special fiber $(\mathcal{Y}_D)_s$ is described by an equation of the form $y^2 = \overline{\rho_0}(x_{\alpha,\beta})$ over the k-line $(\mathcal{X}_D)_s$, where $\overline{\rho_0}$ is, in this case, a normalized reduction of $\rho_{\alpha,\beta}$, and it is not a square.

We introduce the following notation.

Definition 4.33. Given a point $P \in (\mathcal{X}_D)_s$, we define $\mu(\mathcal{X}_D, P)$ to be the order of vanishing of the derivative $\overline{\rho_0}'$ of $\overline{\rho_0}$ at P; when $P = \infty$, we set $\mu(\mathcal{X}_D, P) = 2g - \deg(\overline{\rho_0}')$.

Remark 4.34. We make note of the following.

(a) The integer $\mu(\mathcal{X}_D, P)$ is independent of the chosen good part-square decomposition for $f_{\alpha,\beta}$, thanks to Proposition 4.18(b).
(b) Since $p = 2$, the derivative $\overline{\rho_0}'$ is a square, and so the integer $\mu(\mathcal{X}_D, P)$ is even and non-negative for all $P \in (\mathcal{X}_D)_s$.
(c) Since the degree of ρ is $2g+1$, we have $\sum_{P \in (\mathcal{X}_D)_s} \mu(\mathcal{X}_D, P) = 2g$.

It is immediate to verify that the singularities of $(\mathcal{Y}_D)_s$ lie exactly over the finite set of points $R_{\text{sing}} \subseteq (\mathcal{X}_D)_s$ at which $\overline{\rho_0}'$ vanishes, i.e. the points at which $\mu(\mathcal{X}_D, P) > 0$. Since we have $\sum_{P \in (\mathcal{X}_D)_s} \mu(\mathcal{X}_D, P) = 2g$, and since the integer $\mu(\mathcal{X}_D, P)$ is always even, we have that R_{sing} has cardinality $\leq g$.

We remark that, if $t = 0$, then the points of R_{sing} are just the roots of some (any) normalized reduction of $(f_{\alpha,\beta})'$, because, in this case, the trivial part-square decomposition $f_{\alpha,\beta} = 0^2 + f_{\alpha,\beta}$ is good; in particular, when $t = 0$ we have $R_{\text{mult}} \subseteq R_{\text{sing}}$, where R_{mult} is the set of points of $(\mathcal{X}_D)_s$ to which two or more of the roots $\mathcal{R} \cup \{\infty\}$ reduce.

The normalization $\widetilde{(\mathcal{Y}_D)_s}$ of the special fiber $(\mathcal{Y}_D)_s$ is simply a projective line, and $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ is the Frobenius cover of the projective line $(\mathcal{X}_D)_s$, which can be described by an equation of the form $y^2 = x_{\alpha,\beta}$.

Proposition 4.35. If D is a disc such that $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable, then we have $\mathcal{X}_D \leq \mathcal{X}^{(\text{rst})}$ if and only if $|R_{\text{sing}}| \geq 3$. Moreover, whenever $\mathcal{X} \leq \mathcal{X}^{(\text{rst})}$, the strict transform of $(\mathcal{Y}_D)_s$ in $(\mathcal{Y}^{\text{rst}})_s$ is a projective line.

Proof. The special fiber $(\mathcal{Y}_D)_s$ is reduced, and it consists of a unique component V, which has $|R_{\text{sing}}|$ unibranch singularities lying over $|R_{\text{sing}}|$ distinct points of $(\mathcal{X}_D)_s$; moreover, the normalization $\widetilde{V}$ is a line. We thus have $m(V) = 1$, $a(V) = 0$, $w(V) = |R_{\text{sing}}|$, and $\underline{w}(V) = (1, \ldots, 1)$ and consequently deduce, via the criterion expressed in Proposition 3.17, that $\mathcal{X}_D \leq \mathcal{X}^{(\text{rst})}$ if and only if $N \geq 3$.

Moreover we have that, when $\mathcal{X}_D \leq \mathcal{X}^{(\text{rst})}$, the strict transform of $(\mathcal{Y}_D)_s$ in $(\mathcal{Y}^{\text{rst}})_s$ coincides with the normalization $\widetilde{(\mathcal{Y}_D)_s}$: the proof is identical to the one given in Theorem 4.32, and in our specific case, $\widetilde{(\mathcal{Y}_D)_s} = \widetilde{V}$ is just a projective line. □

Remark 4.36. Since we always have $|R_{\text{sing}}| \leq g$ as shown in the above discussion, when $g = 1$ or $g = 2$, the hypothesis of Proposition 4.35 is never satisfied, hence we never have $\mathcal{X}_D \leq \mathcal{X}^{(\text{rst})}$ if $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable.

Regarding the contribution of $\mathcal{X}_D$ to $\mathcal{X}^{(\text{rst})}$ when $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable, we have the following result.

Proposition 4.37. Letting $\mathfrak{D}$ be the collection of discs corresponding to the model $\mathcal{X}^{(\mathrm{rst})}$ (see §4.2), let us write $\mathfrak{D} = \mathfrak{D}_{\mathrm{sep}} \sqcup \mathfrak{D}_{\mathrm{insep}}$, where $D \in \mathfrak{D}$ belongs to $\mathfrak{D}_{\mathrm{sep}}$ (resp. $\mathfrak{D}_{\mathrm{insep}}$) if the covering map $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is (resp. is not) separable. Then the set $\mathfrak{D}$ can be reconstructed as $(\mathfrak{D}_{\mathrm{sep}})^{\mathrm{ss}}$ following the algorithm presented in §4.2.

Proof. For a disc $D \in \mathfrak{D}_{\mathrm{insep}}$, Proposition 4.35 ensures that $|R_{\mathrm{sing}}| \geq 3$. since Corollary 3.20 says that $R_{\mathrm{sing}} = \mathrm{Ctr}(\mathcal{X}_D, \mathcal{X}^{(\mathrm{rst})})$, we deduce that $|\,\mathrm{Ctr}(\mathcal{X}_D, \mathcal{X}^{(\mathrm{rst})})| \geq 3$. Now the proposition follows from Remark 4.12. $\square$

Roughly speaking, we can conclude that the role of the inseparable components in $(\mathcal{Y}^{\mathrm{rst}})_s$ is inessential: they are just lines that, in light of the proposition above, only get added whenever it is necessary to create room between three or more separable components that would otherwise intersect at the same point and violate semistability. We can consequently focus our attention on the *separable* components of $(\mathcal{Y}^{\mathrm{rst}})_s$, which is to say on the discs D such that $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ for which $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ in a separable cover. Starting in the next section, we will refer to them by the term *valid discs*, and their determination, by Proposition 4.37, suffices to compute the whole $\mathcal{X}^{(\mathrm{rst})}$.

CHAPTER 5

Clusters and valid discs

We begin this section by defining, in §5.1, *clusters (of roots)*, *depths*, and *relative depths* of clusters, and the *cluster picture* associated to the odd-degree polynomial $f(x)$ defining the hyperelliptic curve $Y : y^2 = f(x)$. This notion of "cluster" is essentially the one found in [7, Definition 1.1], although our definition of it varies slightly from the one found there. It is known (see for instance [7, Theorem 1.10]) that the cluster picture completely determines the structure of the special fiber $(\mathcal{Y}^{\min})_s$ of the minimal regular model as long as we are in the $p \neq 2$ setting. Similarly, when $p \neq 2$, a minor variant of this result says that the structure of $(\mathcal{Y}^{\mathrm{rst}})_s$ is also determined entirely by the cluster picture associated to f, in such a way that each component of $(\mathcal{X}^{(\mathrm{rst})})_s$ corresponds to a non-singleton cluster (see Theorem 5.12). In the $p = 2$ setting, however, it is no longer the case that the cluster picture associated to a polynomial f determines the structure of $(\mathcal{Y}^{\min})_s$ or $(\mathcal{Y}^{\mathrm{rst}})_s$. In light of this, in §5.2 we set up the notion of *valid discs* associated to f, so that each one corresponds to a component of $(\mathcal{X}^{(\mathrm{rst})})_s$, and we explore the relationship between these valid discs and clusters associated to f (Theorem 5.13).

1. Clusters

We want to define the *cluster picture* associated to the set $\mathcal{R} \subseteq \bar{K}$ consisting of the $2g + 1$ roots of the polynomial $f(x)$ defining the hyperelliptic curve Y. First, let us introduce a number of invariants attached to a subset $\mathfrak{s} \subseteq \mathcal{R}$.

Definition 5.1. Given a subset $\mathfrak{s} \subseteq \mathcal{R}$, we set

$$d_+(\mathfrak{s}) = \min_{\zeta, \zeta' \in \mathfrak{s}} v(\zeta - \zeta') \in \mathbb{Q} \cup \{+\infty\}; \qquad d_-(\mathfrak{s}) = \max_{\zeta \in \mathfrak{s},\ \zeta' \in \mathcal{R} \smallsetminus \mathfrak{s}} v(\zeta - \zeta') \in \mathbb{Q} \cup \{-\infty\},$$

where we follow the convention that $\min \varnothing = +\infty$ and $\max \varnothing = -\infty$. The number $d_+(\mathfrak{s})$ is named the *(absolute) depth* of $\mathfrak{s}$, while $\delta(\mathfrak{s}) := d_+(\mathfrak{s}) - d_-(\mathfrak{s}) \in \mathbb{Q} \cup \{+\infty\}$ will be also referred to as the *relative depth* of $\mathfrak{s}$. We will use the notation $I(\mathfrak{s})$ to mean the closed interval $[d_-(\mathfrak{s}), d_+(\mathfrak{s})]$, with the convention that $I(\mathfrak{s}) = \varnothing$ whenever $d_+(\mathfrak{s}) < d_-(\mathfrak{s})$.

We are now ready to define the notion of a cluster.

Definition 5.2. Given $\mathfrak{s} \subseteq \mathcal{R}$, we say that $\mathfrak{s}$ is a *cluster* whenever $\mathfrak{s} \neq \varnothing$ and $\delta(\mathfrak{s}) > 0$. The set of pairs $(\mathfrak{s}, d_+(\mathfrak{s}))$, where $\mathfrak{s}$ varies among all clusters of $\mathcal{R}$, is called the *cluster picture* of $\mathcal{R}$.

Remark 5.3. We note the following.

(a) It is elementary to verify that, given a non-empty subset $\mathfrak{s} \subseteq \mathcal{R}$ is a cluster if and only if there exists a disc $D \subseteq \bar{K}$ such that $D \cap \mathcal{R} = \mathfrak{s}$.
(b) We have that $\mathcal{R}$ itself is always a cluster, with $d_-(\mathcal{R}) = -\infty$, $d_+(\mathcal{R})$ finite, and $\delta(\mathcal{R}) = +\infty$.

(c) For every $a \in \mathfrak{s}$, the singleton $\{a\}$ is always a cluster, with $d_-(\mathcal{R})$ finite, $\delta_+(\{s\}) = +\infty$, and $\delta(\{s\}) = +\infty$.

Definition 5.4. For every cluster $\mathfrak{s} \subsetneq \mathcal{R}$, the *parent cluster* of $\mathfrak{s}$ is the smallest cluster $\mathfrak{s}'$ properly containing it; in this situation, we say that $\mathfrak{s}$ is a *child cluster* of $\mathfrak{s}'$. Two distinct clusters having the same parent are said to be *sibling clusters*.

With the notation of the above definition, it is immediate to verify that $d_+(\mathfrak{s}') = d_-(\mathfrak{s})$.

Proposition 5.5. Given a subset $\mathfrak{s} \subseteq \mathcal{R}$, we have the following:

(a) we have $\delta(\mathfrak{s}) > 0$ if and only if $\mathfrak{s}$ is either the empty set or a cluster; and
(b) we have $\delta(\mathfrak{s}) \geq 0$ if and only if $\mathfrak{s}$ is a (possibly empty) union of sibling clusters.

Proof. This follows immediately from definitions. □

The term *cluster picture* is inspired by the fact that the data of a cluster picture can easily be expressed visually. To do so, we represent elements of $\mathcal{R}$ as points and represent proper clusters as loops surrounding the corresponding subsets of points with numbers next to the loops indicating the corresponding depths.

Example 5.6. The cluster picture associated to $\mathcal{R} := \{0, \pi^4, \pi^3, \pi, \pi(1-\pi^4)\}$, where $\pi \in K$ is an element such that $v(\pi) = 1$, may be visualized using the below diagram, in which the all clusters (except the singleton ones) are displayed together with their relative depths (for the cluster $\mathcal{R}$, the label indicates the absolute depth).

cluster picture of $\mathcal{R}$:

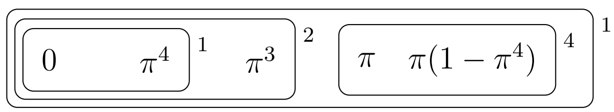

Remark 5.7. Translating a subset $\mathcal{R} \subset \bar{K}$ by an element $\alpha \in \bar{K}$ clearly does not affect the cluster picture. An important automorphism of the projective line is the reciprocal map which takes a finite point $z \neq 0$ to z^{-1} and exchanges 0 and ∞; composed with the translation-by-α map $z \mapsto z - \alpha$, we get an automorphism of the projective line given by $i_\alpha : z \mapsto (z-\alpha)^{-1}$.

One can check readily (see also the proof of [7, Proposition 14.6]) that the cluster picture of a set $\mathcal{R} \subseteq \bar{K}$ transforms in a predictable and easily describable way under such a map i_α. In fact, assume that $\alpha \in \mathcal{R}$, and, for any subset $\mathfrak{s} \subseteq \mathcal{R}$, let $\mathfrak{s}^{\vee,\alpha}$ be defined as:

$$\mathfrak{s}^{\vee,\alpha} = \begin{cases} i_\alpha(\mathcal{R} \smallsetminus \mathfrak{s}) \cup \{0\}, & \text{if } \alpha \in \mathfrak{s}; \\ i_\alpha(\mathfrak{s}), & \text{if } \alpha \notin \mathfrak{s}. \end{cases}$$

This defines a bijection between the subsets of the set of roots of f and the subsets of the set of roots of $f^{\vee,\alpha}$, where $f^{\vee,\alpha}$ is the degree-$(2g+1)$ polynomial defined as $f^{\vee,\alpha} := (z-\alpha)^{2g+2} f((z-\alpha)^{-1})$. Moreover, one readily computes that, when $\alpha \in \mathfrak{s}$, we have $d_\pm(\mathfrak{s}^{\vee,\alpha}) = \mp d_\pm(\mathfrak{s})$, whereas when $\alpha \notin \mathfrak{s}$, we have $d_\pm(\mathfrak{s}^{\vee,\alpha}) = d_\pm(\mathfrak{s}) - 2d_+(\mathfrak{s} \cup \{\alpha\})$. In both cases, we get $\delta(\mathfrak{s}) = \delta(\mathfrak{s}^{\vee,\alpha})$; in particular, we have that $\mathfrak{s}$ is a cluster for f if and only if $\mathfrak{s}^{\vee,\alpha}$ is a cluster for $f^{\vee,\alpha}$.

We want now to introduce some further definitions for later use that relate the clusters to the discs that cut them out of $\mathcal{R}$. We begin with the following remark.

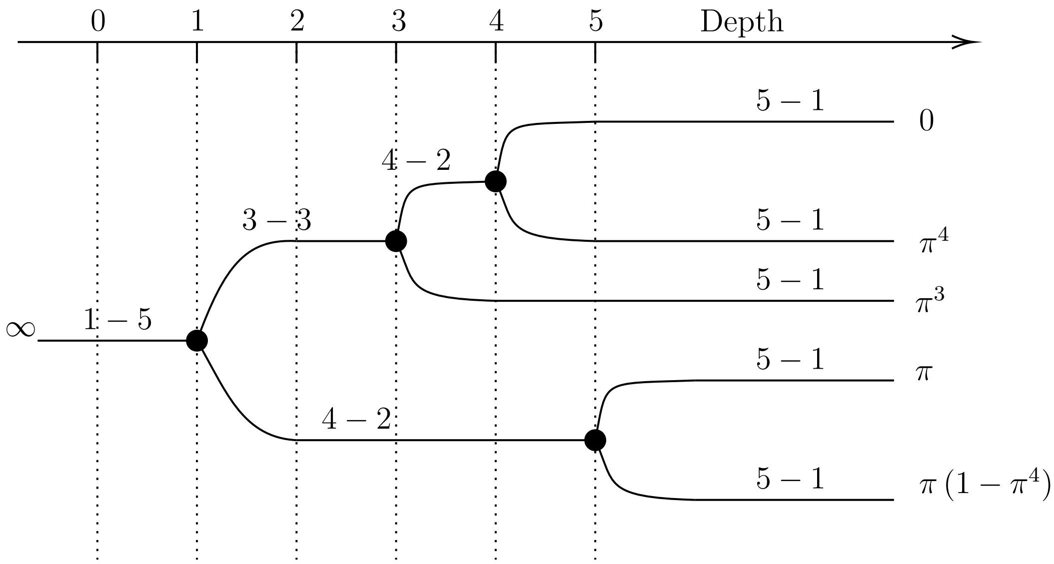

FIGURE 1. This tree describes, when $\mathcal{R}$ is the set of roots of Example 5.6, all discs linked to at least one cluster of $\mathcal{R}$, i.e. all discs D such that $D \cap \mathcal{R} \neq \varnothing$. Each edge corresponds to a cluster $\mathfrak{s} \subseteq \mathcal{R}$ (the labels denote the cardinalities $2g+2-|\mathfrak{s}|$ and $|\mathfrak{s}|$); the initial and final depth of the edge correspond to $d_-(\mathfrak{s})$ and $d_+(\mathfrak{s})$ respectively; the points in the edge correspond to all discs linked to $\mathfrak{s}$. The 4 vertices correspond to those discs that are linked to more than one cluster of $\mathcal{R}$.

Definition 5.8. Given a cluster $\mathfrak{s} \subseteq \mathcal{R}$, we say that a disc $D \subseteq \bar{K}$ is *linked to* $\mathfrak{s} \subseteq \mathcal{R}$ if we have $D = D_{\mathfrak{s},b}$ for some $b \in I(\mathfrak{s})$ (where $D_{\mathfrak{s},b}$ denotes the disc of depth b centered at any point of $\mathfrak{s}$).

To clarify this notion, Figure 1 illustrates all discs linked to each cluster when $\mathcal{R}$ is the cardinality-5 set of roots described in Example 5.6.

Remark 5.9. A disc D is linked to a cluster $\mathfrak{s}$ if and only if we have either $D \cap \mathcal{R} = \mathfrak{s}$ or that D is the minimal disc such that $D \cap \mathcal{R} \supsetneq \mathfrak{s}$. More precisely, if $D = D_{\mathfrak{s},b}$ for some $b \in I(\mathfrak{s})$, we have that $D \cap \mathcal{R} = \mathfrak{s}$ whenever $b \in (d_-(\mathfrak{s}), d_+(\mathfrak{s})]$, whereas D is the smallest disc such that $D \cap \mathcal{R} \supsetneq \mathfrak{s}$ when $b = d_-(\mathfrak{s})$; moreover, in this case the subset $D \cap \mathcal{R} \subseteq \mathcal{R}$ is the parent cluster of $\mathfrak{s}$.

We observe that, given a disc D, the points of $\mathcal{R} \cup \{\infty\}$ reduce to N distinct points $P_1, \ldots, P_{N-1}, \infty \in (\mathcal{X}_D)_s(k)$; we can accordingly write $\mathcal{R} = \mathfrak{s}_1 \sqcup \ldots \sqcup \mathfrak{s}_{N-1} \sqcup \mathfrak{s}_\infty$, where $\mathfrak{s}_i$ consists of the roots of $\mathcal{R}$ reducing to P_i, and $\mathfrak{s}_\infty$ (which is possibly empty) consists of the roots of $\mathcal{R}$ reducing to ∞. We clearly have $D \cap \mathcal{R} = \mathfrak{s}_1 \sqcup \ldots \sqcup \mathfrak{s}_{N-1} = \mathcal{R} \smallsetminus \mathfrak{s}_\infty$. With Figure 1 in mind, it is easy to verify that the following holds.

Lemma 5.10. With notation as above, we have the following.

(a) If $N = 1$, which is equivalent to $D \cap \mathcal{R} = \varnothing$, the disc D is not linked to any cluster.
(b) If $N = 2$, the disc D is linked to exactly one cluster, namely $\mathfrak{s}_1 = \mathcal{R} \smallsetminus \mathfrak{s}_\infty = D \cap \mathcal{R}$.

(c) If $N \geq 3$, the disc D is linked to exactly N clusters, namely $\mathfrak{s}_1, \ldots, \mathfrak{s}_{N-1}$, and $\mathfrak{s}_N := D \cap \mathcal{R} = \mathcal{R} \smallsetminus \mathfrak{s}_\infty = \mathfrak{s}_1 \sqcup \ldots \sqcup \mathfrak{s}_{N-1}$.

Moreover, case (c) occurs if and only if we have $D = D_{\mathfrak{s},d_+(\mathfrak{s})}$ for some non-singleton cluster $\mathfrak{s} \subseteq \mathcal{R}$, in which case we have $\mathfrak{s}_N = \mathfrak{s}$ and that $\mathfrak{s}_1, \ldots, \mathfrak{s}_{N-1}$ are precisely the children clusters of $\mathfrak{s}$, and we have $D = D_{\mathfrak{s}_i,d_-(\mathfrak{s}_i)}$ for $i = 1, \ldots, N-1$.

2. Valid discs

We now define a term which we will use throughout the rest of the work in order to refer to components of the relatively stable model of a hyperelliptic curve.

Definition 5.11. A disc $D \subseteq \bar{K}$ is a *valid disc* if it satisfies $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ and if the quadratic cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is separable.

We note that our notion of *valid disc* differs from the one in [7], although in both cases valid discs are used to build a particular semistable model of Y with desired properties.

The cluster picture allows us to completely identify the valid discs in the tame case.

THEOREM 5.12. In the $p \neq 2$ setting, there is a one-to-one correspondence between non-singleton clusters and clusters discs, which is given by $\mathfrak{s} \mapsto D_{\mathfrak{s},d_+(\mathfrak{s})}$. In other words, the valid discs are precisely those discs that minimally cut out the clusters in $\mathcal{R}$.

PROOF. We have already observed (Lemma 5.10) that a disc D is of the form $D_{\mathfrak{s},d_+(\mathfrak{s})}$ for some non-singleton cluster $\mathfrak{s}$ if and only if the number N of points to which $\mathcal{R} \cup \{\infty\} \subseteq X(K)$ reduces in $(\mathcal{X}_D)_s$ is ≥ 3. Hence, the theorem immediately follows from Theorem 4.32. $\square$

We want some analog of Theorem 5.12 for working over residue characteristic 2; however, in the $p = 2$ setting, valid discs do not correspond in this way in a one-on-one manner with clusters, as is shown by the following theorem.

THEOREM 5.13. In the $p = 2$ setting, we have the following.

(a) Given an odd-cardinality cluster $\mathfrak{s} \subseteq \mathcal{R}$, there is no valid disc D linked to $\mathfrak{s}$.
(b) Given an even-cardinality cluster $\mathfrak{s} \subset \mathcal{R}$, there may be 0, 1, or 2 distinct valid discs D linked to $\mathfrak{s}$.

Remark 5.14. In the $p = 2$ setting, we will see that it is possible that a valid discs D is linked to no cluster, which is to say $D \cap \mathcal{R} = \varnothing$ (see Lemma 5.10), whereas Theorem 5.12 ensures that this never happens when $p \neq 2$.

Remark 5.15. It follows directly from Theorem 5.12 and Lemma 5.10 that the statement in part (b) of the above theorem holds in the $p \neq 2$ setting as well. In fact, when $p \neq 2$, that statement holds even after removing the hypothesis that $\mathfrak{s}$ has even cardinality, and it can be made more precise by saying that there there are *exactly* 2 distinct valid discs D linked to $\mathfrak{s}$, namely $D_{\mathfrak{s},d_+(\mathfrak{s})}$ and $D_{\mathfrak{s},d_-(\mathfrak{s})}$, except in the case $\mathfrak{s} = \mathcal{R}$, when there is exactly 1, namely $D_{\mathfrak{s},d_+(\mathfrak{s})}$.

The proof of Theorem 5.13 is deferred to the following section, in which we set up a framework for considering the models $\mathcal{Y}_D$ corresponding to families of discs $D := D_{\alpha,b}$ which share a common center α and which all contain the same subset $\mathfrak{s} \subseteq \mathcal{R}$ of roots.

CHAPTER 6

Finding valid discs with a given center

In this section, we will fix a center $\alpha \in \bar{K}$, and we will investigate for which depths $b \in I \subseteq \mathbb{Q}$, where $I = [d_-, d_+]$ is some closed interval, the disc $D_{\alpha,b}$ is valid for the hyperelliptic curve $Y : y^2 = f(x)$. The interval I will be chosen so that, when b ranges in the internal of I, the intersection $\mathfrak{s} := D_{\alpha,b} \cap \mathcal{R}$ is constant. More precisely, there are two scenarios we are interested in.

(a) We choose $\mathfrak{s}$ to be a cluster of $\mathcal{R}$, we fix $\alpha \in \bar{K}$ to be any point of the disc $D_{\mathfrak{s},d_+(\mathfrak{s})}$, and we set $I = I(\mathfrak{s}) = [d_-(\mathfrak{s}), d_+(\mathfrak{s})]$. In this case, as we vary $b \in I$, we have that $D_{\alpha,b}$ ranges among all discs linked to $\mathfrak{s}$.

(b) We choose $\mathfrak{s} = \varnothing$, we fix $\alpha \in \bar{K} \smallsetminus \mathcal{R}$, and we let $I = [d_-, +\infty)$, where d_- is the maximum depth such that $D_{\alpha,d_-} \cap \mathcal{R} \neq \varnothing$. This means that, as b ranges in the interior of I, we have that $D_{\alpha,b}$ ranges among all discs centered at α that are linked to no cluster of $\mathcal{R}$.

It is important to consider case (b) as well as case (a), because, when $p = 2$, there may exist valid discs that are linked to no cluster in $\mathcal{R}$.

The section is organized as follows. In §6.1 we introduce the language of translated and scaled part-square decompositions, which will be useful for dealing with the problem. In §6.2 we identify the depths $b \in I$ for which $D_{\alpha,b}$ is a valid disc as the endpoints $b_\pm$ of a sub-interval $J \subseteq I$ (see Theorem 6.18). When $p \neq 2$, we always have $J = I$; however, when $p = 2$, we have that J may be strictly smaller than I; for example, we will see that $J = \varnothing$ whenever $|\mathfrak{s}|$ is odd. Subsections §6.3 and §6.4 develop our strategy for determining J when $\mathfrak{s}$ has even cardinality, provided that, for each of the two factors $f^{\mathfrak{s}}$ and $f^{\mathcal{R} \smallsetminus \mathfrak{s}}$ of f corresponding to the roots lying in $\mathfrak{s}$ and $\mathcal{R} \smallsetminus \mathfrak{s}$ respectively, we know a part-square decomposition that is totally odd with respect to the center α. Finally, in §6.5 we show that, when $\mathfrak{s}$ has even cardinality, the necessary computations can also be performed by replacing totally odd part-square decompositions with *sufficiently odd* ones, which are in general easier to find. In §6.6, we present an algorithm to compute sufficiently odd decompositions, while in §6.7 we present through elementary computations its application to low-degree polynomials.

1. Translated and scaled part-square decompositions

Given any polynomial $h(z) \in \bar{K}[z]$ and any two elements $\alpha \in \bar{K}, \beta \in \bar{K}^\times$, we can compute the (Gauss) valuations of the polynomial $h_{\alpha,\beta}$ obtained from h by translating by α and scaling by β (see §1.5.5.4 for this notational convention). The following lemma will allow us to treat the Gauss valuation of a certain translation and scaling of h as a function on discs.

Lemma 6.1. As we vary $\alpha \in \bar{K}$ and $\beta \in \bar{K}^\times$, the valuation $v(h_{\alpha,\beta})$ depends only on the disc $D := D_{\alpha,v(\beta)}$.

PROOF. Let $\alpha, \alpha' \in \bar{K}$ and $\beta, \beta' \in \bar{K}^\times$ be such that $D_{\alpha,v(\beta)} = D_{\alpha',v(\beta')}$; we must show that $v(h_{\alpha,\beta}) = v(h_{\alpha',\beta'})$. It is clearly sufficient to prove the result when $\alpha = 0$ and $\beta = 1$, in which case $h_{\alpha,\beta} = h$, the assumption $D_{\alpha,v(\beta)} = D_{\alpha',v(\beta')}$ means that $v(\beta') = 0$ and $v(\alpha') \geq 0$, and the claim is verified straightforwardly. □

Given a disc D, we will consequently denote $\underline{v}_h(D)$ the valuation of $h_{\alpha,\beta}$ for any α and β such that $D = D_{\alpha,v(\beta)}$. When a center $\alpha \in \bar{K}$ is fixed, we may consider the function $b \mapsto \underline{v}_h(D_{\alpha,b}) \in \mathbb{Q} \cup \{+\infty\}$ defined for all $b \in \mathbb{Q}$.

Lemma 6.2. Suppose a center $\alpha \in \bar{K}$ is fixed. With the above set-up, we have the following.

(a) The function $b \mapsto \underline{v}_h(D_{\alpha,b}) \in \mathbb{Q} \cup \{+\infty\}$ satisfies the property of being a continuous, non-decreasing piecewise linear function with integer slopes and whose slopes decrease as the input increases.

(b) For any $b \in \mathbb{Q}$ and $\beta \in \bar{K}^\times$ such that $v(\beta) = b$, the left (resp. right) derivative of the function $c \mapsto \underline{v}_h(D_{\alpha,c}) \in \mathbb{Q} \cup \{+\infty\}$ at $c = b$ coincides with the highest (resp. lowest) degree of the variable $x_{\alpha,\beta}$ appearing in the normalized reduction of $h_{\alpha,\beta}$, i.e. with the number of roots ζ of h in $\bar{K}$ (counted with multiplicity) such that $v(\zeta - \alpha) \geq b$ (resp. $v(\zeta - \alpha) > b$).

PROOF. Write H_i for the z^i-coefficient of $h_{\alpha,1}$, and note that $\beta^i H_i$ is the z^i-coefficient of $h_{\alpha,\beta}$ for any scalar β. Now given any $b \in \mathbb{Q}$ and $\beta \in \bar{K}^\times$ with $v(\beta) = b$, by definition we have

$$\underline{v}_h(D_{\alpha,b}) = \min_{0 \leq i \leq \deg(h)} \{v(\beta^i H_i)\} = \min_{0 \leq i \leq \deg(h)} \{v(H_i) + ib\}. \tag{29}$$

All the properties of the function $b \mapsto \underline{v}_h(D_{\alpha,b})$ stated in the lemma immediately follow from the explicit expression given above. □

Given a part-square decomposition $h = q^2 + \rho$ of a nonzero polynomial h, by translating and scaling we can clearly form part-square decompositions $h_{\alpha,\beta} = q^2_{\alpha,\beta} + \rho_{\alpha,\beta}$ for all $\alpha \in \bar{K}, \beta \in \bar{K}^\times$.

Lemma 6.3. Let $h = q^2 + \rho$ be a part-square decomposition.

(a) The property of the induced part-square decomposition $h_{\alpha,\beta} = q^2_{\alpha,\beta} + \rho_{\alpha,\beta}$ being good or not only depends on the disc $D := D_{\alpha,v(\beta)}$ and not on the particular choices of α and β.

(b) The property of the induced part-square decomposition $h_{\alpha,\beta} = q^2_{\alpha,\beta} + \rho_{\alpha,\beta}$ being totally odd or not only depends on our choice of α and not on β.

PROOF. Part (a) is an immediate consequence of Lemma 6.1, while part (b) is immediate. □

We can consequently make the following definitions, which are the variants of those given in Definition 4.15 relative to the choice of a disc.

Definition 6.4. Let $h = q^2 + \rho$ be a part-square decomposition. We make the following definitions:

(a) the decomposition is *good* at a disc D whenever $h_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ is a good part-square decomposition for some (any) $\alpha \in \bar{K}$, $\beta \in \bar{K}^\times$ such that $D = D_{\alpha,v(\beta)}$; and

(b) the decomposition is *totally odd* with respect to a center $\alpha \in \bar{K}$ if $h_{\alpha,\beta} = q_{\alpha,\beta}^2 + \rho_{\alpha,\beta}$ is a totally odd part-square decomposition for some (any) $\beta \in \bar{K}^\times$.

Remark 6.5. If $h = q^2 + \rho$ is a totally odd part-square decomposition with respect to a center α, then, by Corollary 4.19, it is good at the discs $D_{\alpha,b}$, for all $b \in \mathbb{Q}$.

Recalling the number $t_{q,\rho} := v(\rho) - v(h) \in \mathbb{Q} \cup \{+\infty\}$ from §4.3, we define the related function

$$\underline{t}_{q,\rho} := \underline{v}_\rho - \underline{v}_f$$

so that $\underline{t}_{q,\rho}(D) = t_{q_{\alpha,\beta},\rho_{\alpha,\beta}}$ for any $\alpha \in \bar{K}$, $\beta \in \bar{K}^\times$ such that $D = D_{\alpha,v(\beta)}$. When a center $\alpha \in \bar{K}$ is fixed, we can study the function $b \mapsto \underline{t}_{q,\rho}(D_{\alpha,b}) : \mathbb{Q} \to \mathbb{Q} \cup \{+\infty\}$, which is the difference between two continuous piecewise-linear functions and so is itself a continuous piecewise-linear function. Taking into account Remark 4.17, we can give the following definition.

Definition 6.6. Given a (multi-)set of elements $\mathfrak{s} \subseteq \bar{K}$ and a disc D, we define $\mathfrak{t}^{\mathfrak{s}}(D) \in [0, 2v(2)]$ to be $\min\{\underline{t}_{q,\rho}(D), 2v(2)\}$ for any part-square decomposition $h = q^2 + \rho$ which is good at the disc D, where $h(z) \in \bar{K}$ is any polynomial whose set of roots is $\mathfrak{s}$ (counted with multiplicity).

Remark 6.7. Fix a center α. If $h \in \bar{K}[z]$ is a nonzero polynomial and $\mathfrak{s}$ is its (multi-)set of roots, the knowledge of a part-square decomposition $h = q^2 + \rho$ that is totally odd with respect to the center α makes it possible to compute $\mathfrak{t}^{\mathfrak{s}}(D_{\alpha,b}) \in [0, 2v(2)]$ for all depths $b \in \mathbb{Q}$: this follows immediately from Definition 6.6 together with Remark 6.5.

Proposition 6.8. If we have a disjoint union $\mathfrak{s} = \mathfrak{s}_1 \sqcup \ldots \sqcup \mathfrak{s}_N$, then the following hold:

(a) we have $\mathfrak{t}^{\mathfrak{s}}(D) \geq \min\{\mathfrak{t}^{\mathfrak{s}_1}(D), \ldots, \mathfrak{t}^{\mathfrak{s}_N}(D)\}$ for all D; and

(b) the conclusion of (a) is an equality in the following cases:

  (i) whenever the minimum is attained by a unique $\mathfrak{t}^{\mathfrak{s}_i}(D)$; and

  (ii) if $N = 2$, $D \cap \mathfrak{s}_1 = \varnothing$, and there exists a disc $D' \subsetneq D$ such that $\mathfrak{s}_2 \subseteq D'$.

Proof. For $1 \leq i \leq N$, choose polynomials h_i having $\mathfrak{s}_i$ as their sets of roots, and let $h_i = q_i^2 + \rho_i$ be part-square decompositions that are good at the disc D. Then, by setting $q = \prod_i q_i$, we obtain a part-square decomposition for $h := \prod_i h_i$ satisfying $\underline{t}_{q,\rho}(D) \geq \min_i\{\underline{t}_{q_i,\rho_i}(D)\}$ by Proposition 4.21(a). From this part (a) follows. Similarly, points (i) and (ii) of (b) follow straightforwardly from parts (b) and (c) of Proposition 4.21 respectively. □

2. Identifying the valid discs

Let us now consider again the hyperelliptic curve $Y : y^2 = f(x)$, and let us now fix a center $\alpha \in \bar{K}$. In §5.1, we defined, for each subset $\mathfrak{s} \subseteq \mathcal{R}$, the invariants $d_\pm(\mathfrak{s})$ and $\delta(\mathfrak{s})$, as well as the interval $I(\mathfrak{s})$ (see Definition 5.1). We now aim to give analogous definitions *relative to* the center α.

Definition 6.9. Given a subset $\mathfrak{s} \subseteq \mathcal{R}$ and a center $\alpha \in \bar{K}$, we set

$$d_+(\mathfrak{s}, \alpha) = \min_{a \in \mathfrak{s}} v(a - \alpha) \in \mathbb{Q} \cup \{+\infty\}; \quad d_-(\mathfrak{s}, \alpha) = \max_{a \in \mathcal{R} \smallsetminus \mathfrak{s}} v(a - \alpha) \in \mathbb{Q} \cup \{-\infty\}.$$

We also introduce $\delta(\mathfrak{s}, \alpha) := d_+(\mathfrak{s}, \alpha) - d_-(\mathfrak{s}, \alpha) \in \mathbb{Q} \cup \{+\infty\}$, and we use the notation $I(\mathfrak{s}, \alpha)$ to mean the closed interval $[d_-(\mathfrak{s}, \alpha), d_+(\mathfrak{s}, \alpha)]$, with the convention that $I(\mathfrak{s}, \alpha) = \varnothing$ whenever $d_+(\mathfrak{s}, \alpha) < d_-(\mathfrak{s}, \alpha)$.

Remark 6.10. When $\mathfrak{s}$ is a cluster and $\alpha \in D_{\mathfrak{s}, d_+(\mathfrak{s})}$, we have $d_\pm(\mathfrak{s}, \alpha) = d_\pm(\mathfrak{s})$, and $I(\mathfrak{s}, \alpha) = I(\mathfrak{s})$.

Given $\alpha \in \bar{K}$ and $\mathfrak{s} \subseteq \mathcal{R}$, assuming the interval $I(\mathfrak{s}, \alpha)$ has positive length, our aim is to establish for which $b \in I(\mathfrak{s}, \alpha)$ the disc $D_{\alpha,b}$ is a valid disc. When $p \neq 2$, Theorem 5.12 gives an exhaustive answer; we now find a way to address the general case, in which p is arbitrary, i.e. also possibly equal to 2: we will introduce a (possibly empty) closed sub-interval $J(\mathfrak{s}, \alpha)$, whose endpoints, roughly, will correspond to the depths $b \in I(\mathfrak{s}, \alpha)$ for which $D_{\alpha,b}$ is a valid disc, except possibly when $\mathfrak{s} = \varnothing$ (the precise statement is given in Theorem 6.18).

Let us begin by studying the function $I(\mathfrak{s}, \alpha) \ni b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) \in [0, 2v(2)]$, which enjoys the following properties.

Lemma 6.11. *The function $I(\mathfrak{s}, \alpha) \ni b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ is a continuous piecewise-linear function with decreasing slopes. It is identically zero if $|\mathfrak{s}|$ is odd. On the other hand, when $|\mathfrak{s}|$ is even, its slopes are odd integers ranging from $1 - |\mathfrak{s}|$ to $2g + 1 - |\mathfrak{s}|$, except over the subset of $I(\mathfrak{s}, \alpha)$ where $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = 2v(2)$, over which the slope is zero (if this subset contains an open interval).*

Proof. Choose an interior point $b \in I(\mathfrak{s}, \alpha)$, i.e. $b \in (d_-(\mathfrak{s}, \alpha), d_+(\mathfrak{s}, \alpha))$, and choose $\beta \in \bar{K}^\times$ such that $v(\beta) = b$. Any normalized reduction of $f_{\alpha,\beta}$ is a scalar times $x_{\alpha,\beta}^{|\mathfrak{s}|}$. We deduce from Proposition 4.18 that, if $|\mathfrak{s}|$ is odd, the part-square decomposition $f_{\alpha,\beta} = 0^2 + f_{\alpha,\beta}$ is good and $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = 0$, whereas, when $|\mathfrak{s}|$ is even, this decomposition is not good, and we therefore have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) > 0$. In this case, let us take a part-square decomposition $f = q^2 + \rho$ which is totally odd with respect to the center α, so that $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = \min\{\underline{t}_{q,\rho}(D_{\alpha,b}), 2v(2)\}$. Since $\deg(f) = 2g + 1$, by Definition 4.13 the odd-degree polynomial ρ has degree at most $2g + 1$. Now, $b \mapsto \underline{t}_{q,\rho}(D_{\alpha,b})$ is, by definition, the difference between the functions $b \mapsto \underline{v}_\rho(D_{\alpha,b})$ and $b \mapsto \underline{v}_f(D_{\alpha,b})$; by Lemma 6.1, the former is a piecewise linear function with decreasing odd integer slopes between 1 and $2g + 1$, while the latter is linear with slope $|\mathfrak{s}|$ over $I(\mathfrak{s}, \alpha)$. □

In light of the above lemma, either $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ is always $< 2v(2)$ over $I(\mathfrak{s}, \alpha)$, or else it attains the output $2v(2)$ over some closed sub-interval of $I(\mathfrak{s}, \alpha)$ and is $< 2v(2)$ elsewhere. Let $J(\mathfrak{s}, \alpha) = [b_-(\mathfrak{s}, \alpha), b_+(\mathfrak{s}, \alpha)]$ denote the sub-interval of $I(\mathfrak{s}, \alpha) = [d_-(\mathfrak{s}, \alpha), d_+(\mathfrak{s}, \alpha)]$ over which the output of $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ equals $2v(2)$; in the former case just mentioned, we have $J(\mathfrak{s}, \alpha) = \varnothing$, while in the latter case, the interval will have the form $J(\mathfrak{s}, \alpha) = [b_-(\mathfrak{s}, \alpha), b_+(\mathfrak{s}, \alpha)]$ for some endpoints $b_\pm(\mathfrak{s}, \alpha)$.

Remark 6.12. We make the following immediate observations about the subinterval $J(\mathfrak{s}, \alpha) \subseteq I(\mathfrak{s}, \alpha)$.

(a) In the $p \neq 2$ setting, we have $2v(2) = 0$ and so the subinterval $J(\mathfrak{s}, \alpha) \subseteq I(\mathfrak{s}, \alpha)$ coincides with all of $I(\mathfrak{s}, \alpha)$.

(b) In the $p = 2$ setting, if the cluster $\mathfrak{s}$ has odd cardinality, then we have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = 0$ for all $b \in I(\mathfrak{s}, \alpha)$, and therefore we have $J(\mathfrak{s}, \alpha) = \varnothing$.
(c) If $\mathfrak{s} = \varnothing$ (which can only happen in the $p = 2$ setting), then we have $J(\mathfrak{s}, \alpha) \neq \varnothing$ and $b_+(\mathfrak{s}, \alpha) = +\infty$. In fact, the piecewise linear function $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ has only positive slopes by Lemma 6.11, so that $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ becomes equal to $2v(2)$ as $b \to +\infty$.

By Proposition 4.24, it is clear that, given $D = D_{\alpha,b}$ with $b \in I(\mathfrak{s}, \alpha)$, the cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is separable if and only if $b \in J(\mathfrak{s}, \alpha)$; in particular, for $b \in I(\mathfrak{s}, \alpha)$, the disc $D := D_{\alpha,b}$ can only be valid if $b \in J(\mathfrak{s}, \alpha)$. To establish for which $b \in J(\mathfrak{s}, \alpha)$ the disc $D = D_{\alpha,b}$ is valid, we need the following general lemma, which will allow us to compute the ramification of the cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ above 0 and ∞.

Lemma 6.13. Fix a center $\alpha \in \bar{K}$, and choose $b \in \mathbb{Q}$ such that $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = 2v(2)$, and consider the model $\mathcal{Y}_D$ corresponding to the disc $D := D_{\alpha,b}$. Let $\ell(\mathcal{X}_D, P)$ be the integer defined in Definition 4.26 for any point P of $(\mathcal{X}_D)_s$. Write $\partial^+\mathfrak{t}^{\mathcal{R}}$ (resp. $\partial^-\mathfrak{t}^{\mathcal{R}}$) for the right (resp. left) derivative of the function $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$. Then in the $p = 2$ setting, we have the following.

(a) If $\partial^+\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) \geq 0$, then we have $\ell(\mathcal{X}_D, 0) = 0$.
(b) If $\partial^+\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ is odd and negative, then we have $\ell(\mathcal{X}_D, 0) = 1 - \partial^+\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$.
(c) If $\partial^-\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) \leq 0$, then we have $\ell(\mathcal{X}_D, \infty) = 0$.
(d) If $\partial^-\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ is odd and positive, then we have $\ell(\mathcal{X}_D, \infty) = 1 + \partial^-\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$.

In the $p \neq 2$ setting, we instead have the following.

(e) If $\partial^+\underline{v}_f(D_{\alpha,b})$ is even, then we have $\ell(\mathcal{X}_D, 0) = 0$.
(f) If $\partial^+\underline{v}_f(D_{\alpha,b})$ is odd, then we have $\ell(\mathcal{X}_D, 0) = 1$.
(g) If $\partial^-\underline{v}_f(D_{\alpha,b})$ is even, then we have $\ell(\mathcal{X}_D, \infty) = 0$.
(h) If $\partial^-\underline{v}_f(D_{\alpha,b})$ is odd, then we have $\ell(\mathcal{X}_D, \infty) = 1$.

Proof. This is just a rephrasing of Lemma 4.30 using the language introduced in §6.1. To see this, let us fix a part-square decomposition $f = q^2 + \rho$ that is totally odd with respect to the center α, and let f_0, q_0, and ρ_0 be the polynomials involved in the statement of Lemma 4.30: they are defined as appropriate scalings of $f_{\alpha,\beta}$, $q_{\alpha,\beta}$ and $\rho_{\alpha,\beta}$, for some chosen $\beta \in \bar{K}^\times$ such that $v(\beta) = b$.

When $p \neq 2$, the polynomial $\overline{f_0}$ is a normalized reduction of $f_{\alpha,\beta}$, and it is easy to see that parts (e)–(h) of the lemma follow from Lemma 4.30 once the left and right derivatives of $\underline{v}_f$ at $D_{\alpha,b}$ are interpreted in light of Lemma 6.2.

When $p = 2$, the polynomial $\overline{q_0}$ is a normalized reduction of $q_{\alpha,\beta}$, and either the polynomial $\overline{\rho_0}$ is 0 (when $\underline{t}_{q,\rho}(D_{\alpha,\beta}) > 2v(2)$), or it is a normalized reduction of $\rho_{\alpha,\beta}$ (when $\underline{t}_{q,\rho}(D_{\alpha,\beta}) = 2v(2)$). Now we have $\mathfrak{t}(D_{\alpha,c}) = \min\{\underline{t}(D_{\alpha,c}), 2v(2)\}$ for all $c \in \mathbb{Q}$ (see Remark 6.7); moreover, whenever $\underline{t}_{q,\rho}(D_{\alpha,c}) > 0$ (and hence, in particular, for all c in a neighborhood of b), we can write $\underline{t}_{q,\rho}(D_{\alpha,c}) = \underline{v}_\rho(D_{\alpha,c}) - 2\underline{v}_q(D_{\alpha,c})$, where, in light of Lemma 6.2, the first summand only has odd slopes, while the second summand only has even slopes. Let n_ρ and n_q denote the orders of vanishing of $\overline{\rho_0}$ and $\overline{q_0}$ at $x_{\alpha,\beta} = 0$. The assumption in (a) means that either we have $\underline{t}_{q,\rho}(D_{\alpha,b}) > 2v(2)$, or we have $\underline{t}_{q,\rho}(D_{\alpha,b}) = 2v(2)$ with $\partial^+\underline{t}_{q,\rho}(D_{\alpha,b}) \geq 0$; thanks to Lemma 6.2, this can be translated into saying that $n_\rho \geq 2n_q$, and it is now evident that the conclusion of part (a) follows from Lemma 4.30. A similar reasoning can be followed to prove parts (b)–(d). $\square$

As a first application of the lemma above, we will show that a necessary condition for $D_{\alpha,b}$ to be a valid disc when $b \in I(\mathfrak{s}, \alpha)$ is that b is an endpoint of the sub-interval $J(\mathfrak{s}, \alpha) \subseteq I(\mathfrak{s}, \alpha)$.

Lemma 6.14. Given $b \in I(\mathfrak{s}, \alpha)$ and letting $D = D_{\alpha,b}$, we have the following.

(a) If $b \notin J(\mathfrak{s}, \alpha)$, the cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable; hence, we have in particular that D is not a valid disc.
(b) If b is an interior point of $J(\mathfrak{s}, \alpha)$, then we have $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$, and so D is not a valid disc.

Consequently, $D_{\alpha,b}$ can only be a valid disc if $J(\mathfrak{s}, \alpha) \neq \varnothing$ and if b is an endpoint of $J(\mathfrak{s}, \alpha)$, i.e. $b \in \{b_-(\mathfrak{s}, \alpha), b_+(\mathfrak{s}, \alpha)\}$.

Proof. The statement of (a) is a direct result of Proposition 4.24, as we have already discussed. We therefore set out to prove the statement of (b); we assume that $J(\mathfrak{s}, \alpha) \neq \varnothing$ and $b_-(\mathfrak{s}, \alpha) < b < b_+(\mathfrak{s}, \alpha)$ and let $D = D_{\alpha,b}$. The number N of distinct points of $(\mathcal{X}_D)_s$ to which the roots of $\mathcal{R} \cup \{\infty\}$ reduce is at most 2; this is because, since b does not coincide with an endpoint of the interval $I(\mathfrak{s}, \alpha)$, we have that the $2g + 2$ roots $\mathcal{R} \cup \{\infty\}$ each reduce either to 0 or to ∞ in $(\mathcal{X}_D)_s$. Moreover, since b is an interior point of $J(\mathfrak{s}, \alpha)$, we have $\mathfrak{t}^{\mathcal{R}}(D) = 2v(2)$ and that the left and right derivatives of $b' \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b'})$ at $b' = b$ are both equal to 0; by Lemma 6.13, this implies, in the $p = 2$ setting, that $(\mathcal{Y}_D)_s$ has two branches above $0 \in (\mathcal{X}_D)_s$ and two branches above $\infty \in (\mathcal{X}_D)_s$, and the special fiber $(\mathcal{Y}_D)_s$ consequently consists of two rational components (see Remark 4.25).

Now Theorem 4.32 implies that $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$; this is because we know that $N \leq 2$, and, in the $p = 2$ setting, that the special fiber $(\mathcal{Y}_D)_s$ is not irreducible. $\square$

Remark 6.15. The lemma above, applied to the case in which $\mathfrak{s}$ is a cluster and $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$, provides a proof of Theorem 5.13: in fact, the lemma shows that no more than 2 valid discs can be linked to the same cluster and that no valid disc can be linked to $\mathfrak{s}$ if $J(\mathfrak{s}, \alpha) = \varnothing$. By Remark 6.12, this applies in particular when $p = 2$ and $\mathfrak{s}$ has odd cardinality to show that there is no valid disc linked to $\mathfrak{s}$ in this case.

Now assume that $J(\mathfrak{s}, \alpha) \neq \varnothing$. Among the discs $D_{\alpha,b}$ with $b \in I(\mathfrak{s}, \alpha)$, the only candidate valid discs are those of depths $b_-(\mathfrak{s}, \alpha)$ and $b_+(\mathfrak{s}, \alpha)$, as long as these depths are not $\pm\infty$. Let us write $\lambda_-(\mathfrak{s}, \alpha) = \partial^-\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b_-})$ and $\lambda_+(\mathfrak{s}, \alpha) = -\partial^+\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b_+})$ (where $\partial^\pm\mathfrak{t}^{\mathcal{R}}$ is defined as in Lemma 6.13). The integer $\lambda_-(\mathfrak{s}, \alpha)$ (resp. $\lambda_+(\mathfrak{s}, \alpha)$) is only defined if $J(\mathfrak{s}, \alpha) \neq \varnothing$ and its endpoint $b_-(\mathfrak{s}, \alpha)$ (resp. $b_+(\mathfrak{s}, \alpha)$) does not coincide with $d_-(\mathfrak{s}, \alpha)$ (resp. $d_+(\mathfrak{s}, \alpha)$). In particular, $\lambda_+(\mathfrak{s}, \alpha)$ and $\lambda_-(\mathfrak{s}, \alpha)$ can only be defined if $p = 2$ and $\mathfrak{s}$ has even cardinality (by Remark 6.12); when defined, they are both positive odd integers (by Lemma 6.11); more precisely, we have $\lambda_-(\mathfrak{s}, \alpha) \in \{1, 3, \ldots, 2g + 1 - |\mathfrak{s}|\}$ and $\lambda_+(\mathfrak{s}, \alpha) \in \{1, 3, \ldots |\mathfrak{s}| - 1\}$.

Remark 6.16. When $\mathfrak{s}$ is a cluster and $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$, we have already observed in Remark 6.10 that $I(\mathfrak{s}, \alpha) = I(\mathfrak{s})$. It is also evident that, in this case, $b \in I(\mathfrak{s}, \alpha) \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ does not depend on the particular choice of $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$; hence, the sub-interval $J(\mathfrak{s}, \alpha) = [b_-(\mathfrak{s}, \alpha), b_+(\mathfrak{s}, \alpha)]$ and the slopes $\lambda_\pm(\mathfrak{s}, \alpha)$ are independent of this choice as well. Hence, for

$\mathfrak{s}$ a cluster we may (and often will) use without ambiguity the shorter notation $J(\mathfrak{s})$, $b_\pm(\mathfrak{s})$, and $\lambda_\pm(\mathfrak{s})$ to mean $J(\mathfrak{s},\alpha)$, $b_\pm(\mathfrak{s},\alpha)$, and $\lambda_\pm(\mathfrak{s},\alpha)$ for some (any) choice of $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$.

Proposition 6.17. With the above notation, suppose that $J(\mathfrak{s},\alpha) \neq \varnothing$, and let $D_\pm = D_{\alpha,b_\pm(\mathfrak{s},\alpha)}$. Then we have the following.

(a) Assume that $b_-(\mathfrak{s},\alpha) < b_+(\mathfrak{s},\alpha)$ and $|\mathfrak{s}|$ is even. Then we have $\ell(\mathcal{X}_{D_-}, 0) = \ell(\mathcal{X}_{D_+}; \infty) = 0$.

(b) Assume that $b_-(\mathfrak{s},\alpha) < b_+(\mathfrak{s},\alpha)$ and $|\mathfrak{s}|$ is odd. Then we have $\ell(\mathcal{X}_{D_-}, 0) = \ell(\mathcal{X}_{D_+}; \infty) = 1$.

(c) Assume that $b_-(\mathfrak{s},\alpha) > d_-(\mathfrak{s},\alpha)$. Then we have $\ell(\mathcal{X}_{D_-}, \infty) = 1 + \lambda_-(\mathfrak{s},\alpha)$.

(d) Assume that $b_+(\mathfrak{s},\alpha) < d_+(\mathfrak{s},\alpha)$. Then we have $\ell(\mathcal{X}_{D_+}, 0) = 1 + \lambda_+(\mathfrak{s},\alpha)$.

Proof. This follows immediately from Lemma 6.13, taking into account the properties that we have already discussed of the piecewise-linear function $I(\mathfrak{s},\alpha) \ni b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b})$ and of the linear function $I \ni b \mapsto \underline{v}_f(D_{\alpha,b})$ in our setting. $\square$

We are now ready to state a necessary and sufficient condition for $D_{\alpha,b}$ to be a valid disc when $b \in I(\mathfrak{s},\alpha)$.

Theorem 6.18. Let $\mathfrak{s}$ be a cluster of $\mathcal{R}$, and suppose that α is a point in $D_{\mathfrak{s},d_+(\mathfrak{s})}$ or that $\mathfrak{s} = \varnothing$ and α is any point of $\bar{K} \smallsetminus \mathcal{R}$. Let $D = D_{\alpha,b}$ for some $b \in I(\mathfrak{s},\alpha) = [d_-(\mathfrak{s},\alpha), d_+(\mathfrak{s},\alpha)]$; moreover, when $\mathfrak{s} = \varnothing$, let us assume that b is an interior point of $I(\mathfrak{s},\alpha)$, i.e. that $b > d_-(\mathfrak{s},\alpha)$. In other words, when $\mathfrak{s} \neq \varnothing$ we are assuming that D is any disc linked to $\mathfrak{s}$, while, when $\mathfrak{s} = \varnothing$, the disc D may be any disc that contains α and is linked to no cluster.

(a) If $\mathfrak{s} \neq \varnothing$, then the disc D is valid precisely when b is an endpoint of $J(\mathfrak{s})$. Hence, there exist two (possibly coinciding) valid discs $D_{\alpha,b_-(\mathfrak{s})}$ and $D_{\alpha,b_+(\mathfrak{s})}$ linked to $\mathfrak{s}$ when $J(\mathfrak{s}) \neq \varnothing$, and there does not exist a valid disc linked to $\mathfrak{s}$ when $J(\mathfrak{s}) = \varnothing$.

(b) If $\mathfrak{s} = \varnothing$, in which case we have $J(\varnothing,\alpha) = [d_-(\varnothing,\alpha), +\infty)$, we have that D is a valid disc precisely when b coincides with the left endpoint of $J(\varnothing,\alpha)$ and $\lambda_-(\varnothing,\alpha) \geq 3$. Hence, we have two possibilities:
   (i) when $J(\varnothing,\alpha) = I(\varnothing,\alpha)$, or when $J(\varnothing,\alpha) \subsetneq I(\varnothing,\alpha)$ and $\lambda_-(\varnothing,\alpha) = 1$, there does not exist a valid disc centered at α and linked to no cluster; and
   (ii) when $J(\varnothing,\alpha) \subsetneq I(\varnothing,\alpha)$ and $\lambda_-(\varnothing,\alpha) \geq 3$, there exists exactly 1 valid disc centered at α and linked to no cluster.

Proof. The structure of $J(\mathfrak{s},\alpha)$ in the $\mathfrak{s} = \varnothing$ case is discussed in Remark 6.12(c). Moreover, we have already shown that $D = D_{\alpha,b}$ can only be a valid disc when b is an endpoint of $J(\mathfrak{s},\alpha)$ (see Lemma 6.14). So assume from now on that $J(\mathfrak{s},\alpha) \neq \varnothing$ and that $b \in \{b_-(\mathfrak{s},\alpha), b_+(\mathfrak{s},\alpha)\}$. We remark that, since $\mathfrak{t}^{\mathcal{R}}(D) = 2v(2)$, the cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is separable; to determine whether or not D is a valid disc, we may therefore apply the criterion stated in Theorem 4.32. Let N be the integer defined in that theorem.

Assume that b is an endpoint of $I(\mathfrak{s},\alpha)$. By hypothesis, this is only possible when $\mathfrak{s} \neq \varnothing$, in which case we have $I(\mathfrak{s},\alpha) = I(\mathfrak{s})$, and Lemma 5.10 implies that the roots $\mathcal{R} \cup \{\infty\}$ reduce to ≥ 3 distinct points of $(\mathcal{X}_D)_s$ (i.e., $N \geq 3$), and D is certainly a valid disc by Theorem 4.32.

If, instead, the rational number b is an interior point of $I(\mathfrak{s},\alpha)$, then we are in the $p = 2$ setting; the roots of $\mathfrak{s}$ reduce to $0 \in (\mathcal{X}_D)_s$, while those of $\mathcal{R} \smallsetminus \mathfrak{s}$, together with ∞, reduce to $\infty \in (\mathcal{X}_D)_s$. Assume that $b = b_-(\mathfrak{s}, \alpha)$: the $b = b_+(\mathfrak{s}, \alpha)$ case is analogous and will thus be omitted. We know from Proposition 6.17(b) that $\ell(\mathcal{X}_-, \infty) = 1 + \lambda_-(\mathfrak{s}, \alpha) \geq 2$; in particular, $(\mathcal{Y}_D)_s$ has only one branch above $\infty \in (\mathcal{X}_D)_s$ and is consequently irreducible. If $\mathfrak{s} \neq \varnothing$, then we have $N = 2$ and thus the criterion stated in Theorem 4.32 ensures that $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$. If $\mathfrak{s} = \varnothing$, then we have $N = 1$: all roots of $\mathcal{R} \cup \{\infty\}$ reduce to $\infty \in (\mathcal{X}_D)_s$. In this case, the abelian rank of $(\mathcal{Y}_D)_s$, which is the genus of its normalization, is given by $-1 + \ell(\mathcal{X}_-, \infty)/2 = (\lambda_-(\mathfrak{s}, \alpha) - 1)/2$ by Proposition 4.28. Hence, Theorem 4.32 ensures that, for $\mathfrak{s} = \varnothing$, we have $\mathcal{X}_D \leq \mathcal{X}^{(\mathrm{rst})}$ precisely when $\lambda_-(\mathfrak{s}, \alpha) > 1$. $\square$

3. Separating the roots (for an even-cardinality cluster $\mathfrak{s}$)

Let us fix a center $\alpha \in \bar{K}$, and let $\mathfrak{s} \subseteq \mathcal{R}$ be any even-cardinality subset.

3.1. Factoring f. We write the polynomial $f(x)$ as a product $f(x) = cf^{\mathfrak{s}}(x)f^{\mathcal{R} \smallsetminus \mathfrak{s}}(x)$, where c is the leading coefficient of f and write

$$f^{\mathfrak{s}}(x) = \prod_{a \in \mathfrak{s}} (x - a) \qquad \text{and} \qquad f^{\mathcal{R} \smallsetminus \mathfrak{s}}(x) = \prod_{a \in \mathcal{R} \smallsetminus \mathfrak{s}} (x - a). \tag{30}$$

Now let us define $\mathfrak{t}_+^{\mathfrak{s},\alpha}$ and $\mathfrak{t}_-^{\mathfrak{s},\alpha}$ to be the functions on the domain $[0, +\infty)$ given by

$$\mathfrak{t}_+^{\mathfrak{s},\alpha} : b \mapsto \mathfrak{t}^{\mathfrak{s}}(D_{\alpha, d_+(\mathfrak{s},\alpha) - b}) \qquad \text{and} \qquad \mathfrak{t}_-^{\mathfrak{s},\alpha} : b \mapsto \mathfrak{t}^{\mathcal{R} \smallsetminus \mathfrak{s}}(D_{\alpha, b + d_-(\mathfrak{s},\alpha)}).$$

Essentially, the function $\mathfrak{t}_+^{\mathfrak{s},\alpha}$ is defined by evaluating $\mathfrak{t}^{\mathfrak{s}}$ on discs that are enlargements of $D_{\alpha, d_+(\mathfrak{s},\alpha)}$; all such discs contain $\mathfrak{s}$, and $D_{\alpha, d_+(\mathfrak{s})}$ is the minimal disc centered at α with this property. Symmetrically, the function $\mathfrak{t}_-^{\mathfrak{s},\alpha}$ is defined by evaluating $\mathfrak{t}^{\mathcal{R} \smallsetminus \mathfrak{s}}$ at contractions of $D_{\alpha, d_-(\mathfrak{s},\alpha)}$ around the center α: all such discs are disjoint from $\mathcal{R} \smallsetminus \mathfrak{s}$, except the largest one (i.e. $D_{\alpha, d_-(\mathfrak{s},\alpha)}$), which is the minimal disc centered at α that intersects $\mathcal{R} \smallsetminus \mathfrak{s}$.

Proposition 6.19. Both functions $\mathfrak{t}_\pm^{\mathfrak{s},\alpha}$ are strictly increasing on their domains until they reach $2v(2)$ and become constant. Over the part of the domain where $\mathfrak{t}_+^{\mathfrak{s},\alpha}$ (resp. $\mathfrak{t}_-^{\mathfrak{s},\alpha}$) is not constant, its slopes are decreasing odd integers between 1 and $|\mathfrak{s}| - 1$ (resp. between 1 and $2g + 1 - |\mathfrak{s}|$).

Proof. We will only prove the result for $\mathfrak{t}_+^{\mathfrak{s},\alpha}$, as the proof for $\mathfrak{t}_-^{\mathfrak{s},\alpha}$ is analogous. Choose a part-square decomposition for $f^{\mathfrak{s}} = (q^{\mathfrak{s}})^2 + \rho^{\mathfrak{s}}$ that is totally odd with respect to the center α, so that $\mathfrak{t}_+^{\mathfrak{s},\alpha}(b) = \min\{\underline{t}_{q^{\mathfrak{s}},\rho^{\mathfrak{s}}}(D_{\alpha, b_+(\mathfrak{s},\alpha) - b}), 2v(2)\}$ for all $b \in [0, +\infty)$. Since $\mathfrak{s} \subset D_{\alpha, b_+(\mathfrak{s},\alpha) - b}$, we deduce from Lemma 6.2 that $[0, +\infty) \ni b \mapsto \underline{v}_{f^{\mathfrak{s}}}(D_{\alpha, b_+(\mathfrak{s},\alpha) - b})$ has slope 0; on the other hand, the function $[0, +\infty) \ni b \mapsto \underline{v}_{\rho^{\mathfrak{s}}}(D_{\alpha, b_+(\mathfrak{s},\alpha) - b})$ has odd integer slopes between 1 and $|\mathfrak{s}| - 1$. From this, recalling that $\underline{t}_{q^{\mathfrak{s}},\rho^{\mathfrak{s}}} = \underline{v}_{\rho^{\mathfrak{s}}} - \underline{v}_{f^{\mathfrak{s}}}$ by definition, the proposition follows. $\square$

Remark 6.20. We remark that, when $\mathfrak{s} \neq \varnothing$ and $\alpha \in D_{\mathfrak{s}, d_+(\mathfrak{s})}$, the function $\mathfrak{t}_+^{\mathfrak{s},\alpha}$ does not depend on the particular choice of $\alpha \in D_{\mathfrak{s}, d_+(\mathfrak{s})}$; we will therefore use the notation $\mathfrak{t}_+^{\mathfrak{s}}$ to mean $\mathfrak{t}_+^{\mathfrak{s},\alpha}$ where α is some (any) point of $D_{\mathfrak{s}, d_+(\mathfrak{s})}$. On the other hand, the function $\mathfrak{t}_-^{\mathfrak{s},\alpha}(b)$ is the same for all $\alpha \in D_{\mathfrak{s}, d_+(\mathfrak{s})}$ only when $b \in [0, \delta(\mathfrak{s})] \subseteq [0, +\infty)$: when evaluating at such inputs, we may safely drop the superscript α and simply write $\mathfrak{t}_-^{\mathfrak{s}}$ to mean $\mathfrak{t}_-^{\mathfrak{s},\alpha}$ for any $\alpha \in D_{\mathfrak{s}, d_+(\mathfrak{s})}$.

We now remark that the function $I(\mathfrak{s},\alpha) \ni b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) \in [0, 2v(2)]$ we have studied in the previous subsection can be completely recovered from $\mathfrak{t}_{\pm}^{\mathfrak{s},\alpha}$. In fact, we have the following.

Proposition 6.21. Assume that $I(\mathfrak{s},\alpha)$ has positive length (which is always true, for example, when α and $\mathfrak{s}$ are as in the statement of Theorem 6.18). Then, we have

$$\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = \min\{\mathfrak{t}^{\mathfrak{s}}(D_{\alpha,b}), \mathfrak{t}^{\mathcal{R}\smallsetminus\mathfrak{s}}(D_{\alpha,b})\} = \min\{\mathfrak{t}_{+}^{\mathfrak{s},\alpha}(d_{+}(\mathfrak{s},\alpha)-b), \mathfrak{t}_{-}^{\mathfrak{s},\alpha}(b-d_{-}(\mathfrak{s},\alpha))\} \text{ for all } b \in I(\mathfrak{s},\alpha).$$

Proof. It is clearly enough to prove the result for b an interior point of $I(\mathfrak{s},\alpha)$, which will extend by continuity to the endpoints of $I(\mathfrak{s},\alpha)$. For such an input b, we note that the roots $s \in \mathfrak{s}$ satisfy $v(s-\alpha) < b$, while the roots $s \in \mathcal{R} \smallsetminus \mathfrak{s}$ satisfy $v(s-\alpha) > b$. As a consequence, point (ii) of Proposition 6.8(b) applies. ☐

3.2. A standard form for the two factors. Let us introduce the polynomials

$$f_{+}^{\mathfrak{s},\alpha}(z) := \prod_{a\in\mathfrak{s}}(1-\beta_{d_+}^{-1}(a-\alpha)z) \qquad \text{and} \qquad f_{-}^{\mathfrak{s},\alpha}(z) := \prod_{a\in\mathcal{R}\smallsetminus\mathfrak{s}}(1-\beta_{d_-}(a-\alpha)^{-1}z), \tag{31}$$

where the scalars $\beta_{d_\pm} \in \bar{K}^\times$ are chosen to satisfy $v(\beta_{d_\pm}) = d_{\pm}(\mathfrak{s},\alpha)$. These are just transformed versions of $f^{\mathfrak{s}}$ and $f^{\mathcal{R}\smallsetminus\mathfrak{s}}$, normalized so that their constant terms are 1 and all coefficients are integral. More precisely, we have the conversion formulas $f^{\mathfrak{s}} = \beta_{d_+}^{|\mathfrak{s}|}(f_{+}^{\mathfrak{s},\alpha})^{\vee}(\beta_{d_+}^{-1}(z-\alpha))$ and $f^{\mathcal{R}\smallsetminus\mathfrak{s}} = \big(\prod_{a\in\mathcal{R}\smallsetminus\mathfrak{s}}(\alpha-a)\big)f_{-}^{\mathfrak{s},\alpha}(\beta_{d_-}^{-1}(z-\alpha))$, where $(f_{+}^{\mathfrak{s},\alpha})^{\vee}(z) = z^{|\mathfrak{s}|}f_{+}^{\mathfrak{s},\alpha}(1/z)$. Given a part-square decomposition for $f_{+}^{\mathfrak{s},\alpha}$ and for $f_{-}^{\mathfrak{s},\alpha}$, there is an obvious way of producing one for $f^{\mathfrak{s}}$ and $f^{\mathcal{R}\smallsetminus\mathfrak{s}}$, which in turn induces one for $f = cf^{\mathfrak{s}}f^{\mathcal{R}\smallsetminus\mathfrak{s}}$. More precisely, given two part-square decompositions

$$f_{+}^{\mathfrak{s},\alpha} = q_{+}^2 + \rho_{+}, \qquad f_{-}^{\mathfrak{s},\alpha} = q_{-}^2 + \rho_{-},$$

one obtains the decompositions

$$f^{\mathfrak{s}} = [q^{\mathfrak{s}}]^2 + \rho^{\mathfrak{s}}, \qquad f^{\mathcal{R}\smallsetminus\mathfrak{s}} = [q^{\mathcal{R}\smallsetminus\mathfrak{s}}]^2 + \rho^{\mathcal{R}\smallsetminus\mathfrak{s}}, \qquad f = q^2 + \rho$$

by setting $q^{\mathfrak{s}} = \beta_{d_+}^{|\mathfrak{s}|/2}q_{+}^{\vee}(\beta_{d_+}^{-1}(z-\alpha))$, $q^{\mathcal{R}\smallsetminus\mathfrak{s}} = \sqrt{\prod_{a\in\mathcal{R}\smallsetminus\mathfrak{s}}(\alpha-a)}\,q_{-}(\beta_{d_-}^{-1}(z-\alpha))$, and $q = \sqrt{c}q^{\mathfrak{s}}q^{\mathcal{R}\smallsetminus\mathfrak{s}}$ (after making appropriate choices of square roots), where $q_{+}^{\vee}(z) = z^{|\mathfrak{s}|/2}q_{+}(1/z)$.

Remark 6.22. We have the following.

(a) By construction, we have $\underline{t}_{q^{\mathfrak{s}},\rho^{\mathfrak{s}}}(D_{\alpha,b}) = \underline{t}_{q_+,\rho_+}(d_{+}(\mathfrak{s},\alpha)-b)$ and $\underline{t}_{q^{\mathcal{R}\smallsetminus\mathfrak{s}},\rho^{\mathcal{R}\smallsetminus\mathfrak{s}}}(D_{\alpha,b}) = \underline{t}_{q_-,\rho_-}(b-d_{-}(\mathfrak{s},\alpha))$ for all $b \in \mathbb{Q}$; moreover, the above decomposition of $f^{\mathfrak{s}}$ (resp. $f^{\mathcal{R}\smallsetminus\mathfrak{s}}$) is good at $D_{\alpha,b}$ if and only if the above decomposition of $f_{+}^{\mathfrak{s}}$ (resp. $f_{-}^{\mathfrak{s}}$) is good at $d_{+}(\mathfrak{s},\alpha)-b$ (resp. $b-d_{-}(\mathfrak{s},\alpha)$).

(b) It follows from part (a) above that the introduction of $f_{\pm}^{\mathfrak{s}}$ allows us to reinterpret the function $\mathfrak{t}_{\pm}^{\mathfrak{s},\alpha}$ as $[0,+\infty) \ni b \mapsto \mathfrak{t}^{Z_\pm}(D_{0,b})$, where $Z_\pm$ denotes the set of roots of $f_{\pm}^{\mathfrak{s},\alpha}$. We remark that the translation and homotheties that define $f_{\pm}^{\mathfrak{s},\alpha}$ are chosen so that all elements of $Z_\pm$ have valuation ≤ 0 and some element in each of Z_+ and Z_- has valuation 0.

(c) Part (b) above implies that the knowledge of a totally odd part-square decomposition for $f_{\pm}^{\mathfrak{s},\alpha}$ allows us to compute $\mathfrak{t}_{\pm}^{\mathfrak{s},\alpha}$: this is just Remark 6.7. More precisely, if $f_{\pm}^{\mathfrak{s},\alpha} = q_{\pm}^2 + \rho_{\pm}$ is a totally odd part-square decomposition, we have $\mathfrak{t}_{\pm}^{\mathfrak{s},\alpha}(b) = \min\{\underline{t}_{q_\pm,\rho_\pm}(D_{0,b}), 2v(2)\}$ for all $b \in [0,+\infty)$.

The following proposition will be useful in that it allows one to study the valid discs containing an even-cardinality subset $\mathfrak{s}$ by considering the image of $(\mathcal{R} \cup \{\infty\}) \smallsetminus \mathfrak{s}$ under the reciprocal map (see Remark 5.7 above).

Proposition 6.23. Assume the notation of Remark 5.7; let $\mathfrak{s} \subseteq \mathcal{R}$ a subset of even cardinality; and assume that $\alpha \in \mathfrak{s}$. Then we have $f_\pm^{\mathfrak{s}^\vee,\alpha,0} = f_\mp^{\mathfrak{s},\alpha}$. It follows that, given part-square decompositions $f_\pm^{\mathfrak{s},\alpha} = q_\pm^2 + \rho_\pm$, we have part-square decompositions $f_\pm^{\mathfrak{s}^\vee,\alpha,0} = q_\mp^2 + \rho_\mp$, and in fact, we have the equality of functions $\mathfrak{t}_\pm^{\mathfrak{s}^\vee,\alpha,0} = \mathfrak{t}_\mp^{\mathfrak{s},\alpha}$.

Proof. The first claims can be straightforwardly checked directly from the observations in Remark 5.7 and the defining formulas for the terms. The final claim follows from choosing the decompositions $f_\pm^{\mathfrak{s},\alpha} = q_\pm^2 + \rho_\pm$ to be totally odd and using Remark 6.7. □

3.3. Reconstructing the invariants. Let $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ be the least value of $b \in [0, +\infty)$ at which $\mathfrak{t}_\pm^{\mathfrak{s},\alpha} : [0, +\infty) \to [0, 2v(2)]$ attains $2v(2)$, and let $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ denote the left derivative of $\mathfrak{t}_\pm^{\mathfrak{s},\alpha}$ at $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$, which is clearly only defined when $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha}) > 0$. These invariants are closely related to those introduced in the previous subsection.

Proposition 6.24. Suppose that $\mathfrak{s}$ has even cardinality, that $I(\mathfrak{s}, \alpha)$ has positive length (which always occurs, for example, if $\mathfrak{s}$ and α as as in Theorem 6.18), and that $J(\mathfrak{s}, \alpha) \neq \varnothing$. Then, we have

$$b_\pm(\mathfrak{s}, \alpha) = d_\pm \mp b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha}) \qquad \text{and} \qquad \lambda_\pm(\mathfrak{s}, \alpha) = \lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha}). \tag{32}$$

Proof. We have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = \min\{\mathfrak{t}_+^{\mathfrak{s},\alpha}(b_+ - b), \mathfrak{t}_-^{\mathfrak{s},\alpha}(b - b_-)\}$ by Proposition 6.21 and that each of the functions $\mathfrak{t}_\pm^{\mathfrak{s},\alpha}$ is strictly increasing until it reaches $2v(2)$ by Proposition 6.19. The proposition now follows immediately. □

Remark 6.25. In the context of Proposition 6.24, when $J(\mathfrak{s}, \alpha) = \varnothing$ the formulas in (32) can be taken as the *definitions* of the rational numbers $b_\pm(\mathfrak{s}, \alpha)$ and $\lambda_\pm(\mathfrak{s}, \alpha)$, and from $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) = \min\{\mathfrak{t}_+^{\mathfrak{s},\alpha}(b_+ - b), \mathfrak{t}_-^{\mathfrak{s},\alpha}(b - b_-)\}$, it is easy to see that the condition $J(\mathfrak{s}, \alpha) = \varnothing$ corresponds to the fact that $b_-(\mathfrak{s}, \alpha) > b_+(\mathfrak{s}, \alpha)$: roughly speaking, $J(\mathfrak{s}, \alpha)$ is empty whenever its endpoints, which can always be defined, are in the reversed order. Observe that $\lambda_+(\mathfrak{s}, \alpha) \in \{1, \ldots, |\mathfrak{s}| - 1\}$ is actually only defined when $b_+(\mathfrak{s}, \alpha) > d_+(\mathfrak{s}, \alpha)$ (that is, when $b_0(\mathfrak{t}_+^{\mathfrak{s},\alpha}) > 0$), and $\lambda_-(\mathfrak{s}, \alpha) \in \{1, \ldots, 2g + 1 - |\mathfrak{s}|\}$ is only defined when $d_-(\mathfrak{s}, \alpha) < b_-(\mathfrak{s}, \alpha)$ (that is, when $b_0(\mathfrak{t}_-^{\mathfrak{s},\alpha}) > 0$). When $\mathfrak{s}$ is a cluster and $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$, even with these more general definitions, $\lambda_+(\mathfrak{s}, \alpha)$ and $b_+(\mathfrak{s}, \alpha)$ only depend on $\mathfrak{s}$ and not on the particular choice of the center $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$; the same is true for $b_-(\mathfrak{s}, \alpha)$ and $\lambda_-(\mathfrak{s}, \alpha)$, provided that $b_-(\mathfrak{s}, \alpha) \leq d_+(\mathfrak{s}, \alpha)$ (see Remark 6.20).

The computations of the invariants $J(\mathfrak{s}, \alpha), b_\pm(\mathfrak{s}, \alpha)$ and $\lambda_\pm(\mathfrak{s}, \alpha)$ appearing in Theorem 6.18 are now reduced to determining $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ and $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$. A priori, this would require the knowledge of the functions $\mathfrak{t}_\pm^{\mathfrak{s},\alpha} : [0, +\infty) \to [0, 2v(2)]$, which in turn are immediate to compute once a totally odd part-square decomposition for the polynomials $f_\pm^{\mathfrak{s},\alpha}$ is known: see Remark 6.22(c). However, determining a totally odd part-square decomposition for $f_\pm^{\mathfrak{s},\alpha}$ can be a difficult task, even if easier than determining one for the whole polynomial $f_{\alpha,1}$. In §6.5, we will introduce a class of decompositions which we will name *sufficiently odd decompositions* (see Definition 6.37 below); these are easier to compute and will still allow us to find $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ and $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$, as we will show in Proposition 6.41.

4. Estimating thresholds for depths of even-cardinality clusters

The results that we have obtained in the above subsections show that given an even-cardinality cluster $\mathfrak{s}$ of roots associated to the hyperelliptic curve $Y : y^2 = f(x)$, there are 0, 1, or 2 valid discs linked to it, and the results suggest how we may determine how many valid discs are linked to it via the knowledge of the rational numbers $b_\pm(\mathfrak{s}) = d_\pm(\mathfrak{s}) \mp b_0(\mathfrak{t}_\pm^{\mathfrak{s}})$. Roughly speaking, the results of §6.2 and §6.3 show that an even-cardinality cluster $\mathfrak{s}$ has 2 (resp. 1) valid discs linked to it if and only if its relative depth $\delta(\mathfrak{s}) = d_+(\mathfrak{s}) - d_-(\mathfrak{s})$ exceeds (resp. equals) some threshold depending on $\mathfrak{s}$, namely the rational number given by $b_0(\mathfrak{t}_+^{\mathfrak{s}}) + b_0(\mathfrak{t}_-^{\mathfrak{s}})$. The precise statement is the following rephrasing of Theorem 6.18(a) combined with Remark 6.25.

Proposition 6.26. Given an even-cardinality cluster $\mathfrak{s} \subset \mathcal{R}$ of relative depth $\delta(\mathfrak{s})$ and writing $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_+^{\mathfrak{s}}) + b_0(\mathfrak{t}_-^{\mathfrak{s}})$, there are exactly 2 (resp. 1; resp. 0) valid discs linked to $\mathfrak{s}$ if we have $\delta(\mathfrak{s}) > B_{f,\mathfrak{s}}$ (resp. $\delta(\mathfrak{s}) = B_{f,\mathfrak{s}}$; resp. $\delta(\mathfrak{s}) < B_{f,\mathfrak{s}}$).

Remark 6.27. We note that the rational number $B_{f,\mathfrak{s}}$ given in the above corollary does not depend on the depth $\delta(\mathfrak{s})$ in the following sense. Given a center $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$, let $\mathfrak{s}_{[\lambda]} = \{\lambda(a-\alpha)+\alpha \mid a \in \mathfrak{s}\}$ for some $\lambda \in \bar{K}^\times$ such that $v(\lambda) > -\delta(\mathfrak{s})$, so that $\mathfrak{s}_{[\lambda]}$ is a scaled version of $\mathfrak{s}$ and is a cluster in $\mathcal{R}_{[\lambda]} := \mathfrak{s}_{[\lambda]} \sqcup (\mathcal{R} \smallsetminus \mathfrak{s})$ with relative depth $\delta(\mathfrak{s}_{[\lambda]}) = \delta(\mathfrak{s}) + v(\lambda)$. Then it follows easily from the constructions in §6.3 that we have $\mathfrak{t}_+^{\mathfrak{s}} = \mathfrak{t}_+^{\mathfrak{s}_{[\lambda]}}$ and $\mathfrak{t}_-^{\mathfrak{s}} = \mathfrak{t}_-^{\mathfrak{s}_{[\lambda]}}$, from which it follows that $B_{f_{[\lambda]},\mathfrak{s}_{[\lambda]}} = B_{f,\mathfrak{s}}$. In this sense, loosely speaking, we may view $B_{f,\mathfrak{s}}$ as a sort of "threshold" for the depth of $\mathfrak{s}$ at which we obtain 1 valid disc linked to $\mathfrak{s}$ and above which we obtain 2 valid discs linked to $\mathfrak{s}$.

In the rest of this subsection, we work towards obtaining estimates and exact formulas for the "threshold depth" $B_{f,\mathfrak{s}}$ defined in Proposition 6.26 under various conditions on $\mathfrak{s} \subset \mathcal{R}$.

Proposition 6.28. Let $\mathfrak{s} \subseteq \mathcal{R}$ be an even-cardinality cluster.

(a) Suppose that $\mathfrak{s}$ can be written as a disjoint union $\mathfrak{r} \sqcup \mathfrak{c}_1 \sqcup \ldots \sqcup \mathfrak{c}_N$ for some $N \geq 0$, and where each $\mathfrak{c}_i$ is an even-cardinality child of $\mathfrak{s}$. Let $\delta = d_+(\mathfrak{r}) - d_+(\mathfrak{s})$ (so that, in particular, $\delta = 0$ when $N = 0$, and $\delta = \delta(\mathfrak{r})$ when $N \geq 1$), and assume that $\delta(\mathfrak{c}_i) > \delta$ for all $i = 1, \ldots, N$. Assume moreover that the sum $\sigma := \sum_{a \in \mathfrak{r}}(a - \alpha_0)$ for some (any) fixed $\alpha_0 \in \mathfrak{r}$ has valuation equal to $d_+(\mathfrak{r})$. Then for all $b \in [0, +\infty)$ we have
$$\mathfrak{t}_+^{\mathfrak{s}}(b) = \min\{\delta + b, 2v(2)\},$$
so that $b_0(\mathfrak{t}_+^{\mathfrak{s}}) = \max\{2v(2) - \delta, 0\}$.

(b) Assume that the parent cluster $\mathfrak{c}_1$ of $\mathfrak{s}$ has even cardinality and can be written as a disjoint union $\mathfrak{c}_1 = \mathfrak{c}_2 \sqcup \ldots \sqcup \mathfrak{c}_N \sqcup \mathfrak{s} \sqcup \mathfrak{r}$ for some $N \geq 1$, where $\mathfrak{c}_2, \ldots, \mathfrak{c}_N$ are even-cardinality sibling clusters of $\mathfrak{s}$. Let $\delta = \delta(\mathfrak{r})$, and assume that $\delta(\mathfrak{c}_i) > \delta$ for all $i = 1, \ldots, N$. Assume moreover that the sum $\sigma := \sum_{a \in \mathfrak{r}}(a - \alpha_0)$ has valuation $d_+(\mathfrak{r})$ for some (any) fixed $\alpha_0 \in \mathfrak{r}$. Then for all $b \in [0, \delta(\mathfrak{s})]$ we have
$$\mathfrak{t}_-^{\mathfrak{s}}(b) = \min\{\delta + b, 2v(2)\},$$
so that $b_0(\mathfrak{t}_-^{\mathfrak{s}}) = \max\{2v(2) - \delta, 0\}$.

(c) Let $\mathfrak{r} = \mathfrak{s} \sqcup \mathfrak{c}_1 \sqcup \ldots \sqcup \mathfrak{c}_N$ be a union of $\mathfrak{s}$ and some even-cardinality sibling clusters $\mathfrak{c}_1, \ldots, \mathfrak{c}_N$ of $\mathfrak{s}$, where $N \geq 0$. Let $\delta = d_-(\mathfrak{s}) - d_-(\mathfrak{r})$ (so that, in particular,

$\delta = 0$ when $N = 0$, and $\delta = \delta(\mathfrak{r})$ when $N \geq 1$), and assume that $\delta(\mathfrak{c}_i) > \delta$ for all $i = 1, \ldots, N$. Assume moreover that the sum $\sigma := \sum_{a \in \mathcal{R} \smallsetminus \mathfrak{r}} (a - \alpha_0)^{-1}$ has valuation equal to $-d_-(\mathfrak{r})$ for some (any) fixed $\alpha_0 \in \mathfrak{s}$. Then for all $b \in [0, \delta(\mathfrak{s})]$, we have

$$\mathfrak{t}^{\mathfrak{s}}_-(b) = \min\{\delta + b, 2v(2)\},$$

so that $b_0(\mathfrak{t}^{\mathfrak{s}}_-) = \max\{2v(2) - \delta, 0\}$.

Remark 6.29. The assumption on $v(\sigma)$ in (a) and (b) is automatically satisfied whenever $\mathfrak{r}$ is a disjoint union of two odd-cardinality clusters $\mathfrak{r}_1$ and $\mathfrak{r}_2$ (in particular, it is satisfied whenever $\mathfrak{r}$ has cardinality 2). In fact, fixing choices of elements $\alpha_i \in \mathfrak{r}_i$ for $i = 1, 2$ and choosing $\alpha_0 = \alpha_1$, we may write

$$\sigma = \sum_{a \in \mathfrak{r}} (a - \alpha_1) = |\mathfrak{r}_2|(\alpha_2 - \alpha_1) + \sum_{a \in \mathfrak{r}_1} (a - \alpha_1) + \sum_{a \in \mathfrak{r}_2} (a - \alpha_2).$$

Clearly the valuation of $|\mathfrak{r}_2|(\alpha_2 - \alpha_1)$ equals $d_+(\mathfrak{r})$ while all other terms in the above formula have higher valuation, and therefore the entire sum has valuation $d_+(\mathfrak{r})$.

Analogously, the assumption on $v(\sigma)$ in (c) is automatically satisfied whenever there exist odd-cardinality clusters $\mathfrak{r}_1, \mathfrak{r}_2 \subseteq \mathcal{R}$ such that $\mathfrak{r}_1 = \mathfrak{r} \sqcup \mathfrak{r}_2$; this always happens, in particular, when $\mathfrak{r}$ has cardinality $2g$.

Proof (of Proposition 6.28). Let us first assume that $\mathfrak{s}$ satisfies the hypotheses of part (a). We first set out to show that we have $\mathfrak{t}^{\mathfrak{r}}_+(b) = \min\{b, 2v(2)\}$ for all $b \in [0, +\infty)$ (which in particular proves part (a) in the case that $N = 0$). We consider the polynomial $f_+^{\mathfrak{r},\alpha_0}$ (with the notation in §6.3.3.2): by construction, its coefficient are all integral, and its constant term is 1; moreover, the assumption on $v(\sigma)$ easily implies that its linear coefficient is a unit. From this, it is easy to deduce that ρ_+ is also a polynomial with integral coefficients and unit linear coefficient, where $f_+^{\mathfrak{r},\alpha_0} = q_+^2 + \rho_+$ is a totally odd part-square decomposition. Let us now consider the function $\mathfrak{t}^{\mathfrak{r}}_+ : [0, +\infty) \to \mathbb{Q}$, which is piecewise-linear with odd decreasing slopes between 1 and $|\mathfrak{r}| - 1$ until it reaches $2v(2)$ (see Proposition 6.19). This function can be computed as $\mathfrak{t}^{\mathfrak{r}}_+(b) = \underline{t}_{q_+,\rho_+}(D_{0,b}) = \underline{v}_{\rho_+}(D_{0,b})$; since ρ_+ has integral coefficients and unit linear term, the function $\mathfrak{t}^{\mathfrak{r}}_+$ has initial output 0 and initial slope 1, from which we conclude that $\mathfrak{t}^{\mathfrak{r}}_+(b) = \min\{b, 2v(2)\}$ for all $b \in [0, +\infty)$.

Now fix any $i \in \{1, \ldots, N\}$, and take any disc $D := D_{\mathfrak{s}, d_+(\mathfrak{s}) - b}$ with $b \in [0, +\infty)$. Then the formula for $\mathfrak{t}^{\mathfrak{r}}_+$ that we found above is equivalent to the formula $\mathfrak{t}^{\mathfrak{r}}(D) = \mathfrak{t}^{\mathfrak{r}}_+(b + \delta) = \min\{b + \delta, 2v(2)\}$. Moreover, using the fact that $\mathfrak{t}^{\mathfrak{c}_i}_+$ has positive integer slopes as long as its output is $< 2v(2)$ by Proposition 6.19, we get $\mathfrak{t}^{\mathfrak{c}_i}(D) = \mathfrak{t}^{\mathfrak{c}_i}_+(b + \delta(\mathfrak{c}_i)) \geq \min\{b + \delta(\mathfrak{c}_i), 2v(2)\}$. Now, using our assumption that $\delta(\mathfrak{c}_i) > \delta$, Proposition 6.8(b) implies that $\mathfrak{t}^{\mathfrak{s}}(D) = \mathfrak{t}^{\mathfrak{r}}(D)$, which can clearly be rewritten as $\mathfrak{t}^{\mathfrak{s}}_+(b) = \min\{\delta + b, 2v(2)\}$. This finishes the proof of part (a).

Now, if we apply to the setting described in the hypothesis of (a) the automorphism $i_\alpha : z \mapsto (z - \alpha)^{-1}$ of the projective line, where α is an element of $\mathfrak{c}_1$ (resp. of $\mathfrak{r}$), we obtain exactly the setting described in the hypothesis of (b) (resp. of (c)): this can be readily verified by applying Remark 5.7. Moreover, Proposition 6.23 ensures that, after applying i_α, the conclusion of (a) turns into that of (b) (resp. of (c)). $\square$

Proposition 6.30. We have the following analogous statements.

(a) Suppose that $\mathfrak{s}$ is an even-cardinality cluster which itself is the disjoint union of even-cardinality child clusters $\mathfrak{c}_1, \ldots, \mathfrak{c}_N$ for some $N \geq 2$. The minimum of the set
$$\{\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i))\}_{1 \leq i \leq N} \cup \{\mathfrak{t}_+^{\mathfrak{s}}(0)\}$$
of rational numbers is attained by more than one element. In particular, if we have $\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i)) = 2v(2)$ for $1 \leq i \leq N$, then we have $\mathfrak{t}_+^{\mathfrak{s}}(0) = 2v(2)$ also.

(b) Suppose that $\mathfrak{s}$ is an even-cardinality cluster whose parent cluster $\mathfrak{s}'$ is a disjoint union $\mathfrak{s} \sqcup \mathfrak{c}_2 \sqcup \ldots \sqcup \mathfrak{c}_N$ of even-cardinality clusters for some $N \geq 2$. Then, the minimum of the set
$$\{\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i))\}_{2 \leq i \leq N} \cup \{\mathfrak{t}_-^{\mathfrak{s}'}(\delta(\mathfrak{s}')), \mathfrak{t}_-^{\mathfrak{s}}(0)\}$$
of rational numbers is attained by more than one element. In particular, if we have $\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i)) = 2v(2)$ for $2 \leq i \leq N$, and $\mathfrak{t}_-^{\mathfrak{s}'}(\delta(\mathfrak{s}')) = 2v(2)$ then we have $\mathfrak{t}_-^{\mathfrak{s}}(0) = 2v(2)$ also.

PROOF. Let us assume the setting of (a). If the minimum of the set $\{\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i)\}_{1 \leq i \leq N}$ is attained by more than one element, then we are done, so assume that this minimum is attained by a unique element. Then, if we apply Proposition 6.8 to the disc $D := D_{\mathfrak{s}, d_+(\mathfrak{s})}$, we obtain that $\mathfrak{t}_+^{\mathfrak{s}}(0) = \min_{1 \leq i \leq N} \mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i))$, and the claim is proved. The proof of part (b) is analogous. □

The following pleasant corollary provides a simple result for a very special case of cluster picture.

Corollary 6.31. Suppose that we have a cluster $\mathfrak{S}$ of cardinality $2g$ which has g children $\mathfrak{c}_1, \ldots, \mathfrak{c}_g$, each of which has cardinality 2. Assume that we have $\delta(\mathfrak{c}_i) \geq 2v(2)$ for $1 \leq i \leq g$ as well as $\delta(\mathfrak{S}) \geq 2v(2)$. Then we have $b_0(\mathfrak{t}_+^{\mathfrak{c}_i}) = 2v(2)$ and $b_0(\mathfrak{t}_-^{\mathfrak{c}_i}) = 0$ for $1 \leq i \leq g$, and we have $b_0(\mathfrak{t}_+^{\mathfrak{S}}) = 0$ and $b_0(\mathfrak{t}_-^{\mathfrak{S}}) = 2v(2)$.

PROOF. Letting $\mathfrak{s} = \mathfrak{r} = \mathfrak{c}_i$ and applying Proposition 6.28(a), we get $\mathfrak{t}_+^{\mathfrak{c}_i}(b) = \min\{b, 2v(2)\}$ and $b_0(\mathfrak{t}_+^{\mathfrak{c}_i}) = 2v(2)$ for each i; similarly, if we let $\mathfrak{s} = \mathfrak{r} = \mathfrak{S}$ and apply Proposition 6.28(c), we get $\mathfrak{t}_-^{\mathfrak{S}}(b) = \min\{b, 2v(2)\}$ and $b_0(\mathfrak{t}_-^{\mathfrak{S}}) = 2v(2)$. Now, by applying Proposition 6.30 and using the assumption that $\delta(\mathfrak{c}_i) \geq 2v(2)$ and $\delta(\mathfrak{S}) \geq 2v(2)$, we get $\mathfrak{t}_+^{\mathfrak{S}}(0) = 2v(2)$ and $\mathfrak{t}_-^{\mathfrak{c}_i}(0) = 2v(2)$ for all i, which imply $b_0(\mathfrak{t}_+^{\mathfrak{S}}) = 0$ and $b_0(\mathfrak{t}_-^{\mathfrak{c}_i}) = 0$, respectively. □

Remark 6.32. The special case treated by the above corollary can be viewed more symmetrically as involving a hyperelliptic curve defined by a degree-$(2g + 2)$ polynomial (after applying a suitable automorphism so that the $2g+2$ branch points of Y all have x-coordinate different from ∞) whose roots are paired into $g+1$ clusters each of cardinality 2; this is the type of hyperelliptic curve treated in [8].

It immediately follows from Corollary 6.31 combined with Theorem 6.18(a) and Proposition 6.24 that under the hypotheses of Corollary 6.31, letting $d = d_+(\mathfrak{S}) = d_-(\mathfrak{c}_i)$, the full set of valid discs consists of $D_{\mathfrak{S},d}$, along with $D_{\mathfrak{S},d-\delta(\mathfrak{S})+2v(2)}$ if $\delta(\mathfrak{S}) > 2v(2)$, as well as $D_{\mathfrak{s}_i, d+\delta(\mathfrak{c}_i)-2v(2)}$ for each i such that $\delta(\mathfrak{c}_i) > 2v(2)$. In fact, it is not difficult to see that the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to the valid discs other than $D_{\mathfrak{S},d}$ are all vertical (-2)-curves which, when contracted, yield the stable model (which in turn coincides with $\mathcal{Y}_{D_{\mathfrak{S},d}}$). The special fiber of the stable model then has a node corresponding to each i

such that $\delta(\mathfrak{c}_i) > 2v(2)$ as well as an additional node if $\delta(\mathfrak{S}) > 2v(2)$. This essentially recovers the statement of [8, Proposition 1.5].

Lemma 6.33. Let $\mathfrak{s}$ be a cluster of even cardinality. We have

$$b_0(\mathfrak{t}_+^{\mathfrak{s}}) \geq \frac{2v(2) - \mathfrak{t}_+^{\mathfrak{s}}(0)}{|\mathfrak{s}| - 1} \quad \text{and} \quad b_0(\mathfrak{t}_-^{\mathfrak{s}}) \geq \frac{2v(2) - \mathfrak{t}_-^{\mathfrak{s}}(0)}{2g + 1 - |\mathfrak{s}|}. \tag{33}$$

Moreover, if $\lambda_+(\mathfrak{s}) = |\mathfrak{s}| - 1$ (resp. $\lambda_-(\mathfrak{s}) = 2g + 1 - |\mathfrak{s}|$), then the first (resp. the second) inequality above is an equality.

Proof. This follows immediately from the properties of $\mathfrak{t}_\pm^{\mathfrak{s}}$ presented in Proposition 6.19. $\square$

Lemma 6.34. Let $\mathfrak{s}$ be a cluster of even cardinality. Then we have the following:

(a) if $\mathfrak{s}$ has an odd-cardinality child cluster, then we have $\mathfrak{t}_+^{\mathfrak{s}}(0) = 0$; and
(b) if $\mathfrak{s}$ has an odd-cardinality sibling cluster, then we have $\mathfrak{t}_-^{\mathfrak{s}}(0) = 0$.

Proof. It is immediate to see that, if $\mathfrak{s}$ has a child cluster $\mathfrak{c}$ of odd cardinality $2m+1$, then, letting $\alpha \in \mathfrak{c}$, any normalized reduction of $f_+^{\mathfrak{s},\alpha}$ has odd degree $2g+1-2m$ and thus, in particular, is not a square. This implies that $f_+^{\mathfrak{s},\alpha} = 0^2 + f_+^{\mathfrak{s},\alpha}$ is a good part-square decomposition (see Proposition 4.18), and hence that $\mathfrak{t}_+^{\mathfrak{s}}(0) = 0$. This proves (a); the proof of (b) is analogous. $\square$

Proposition 6.35. Let $\mathfrak{s}$ be a cluster of even cardinality, and let $\mathfrak{s}'$ be its parent cluster. The rational number $B_{f,\mathfrak{s}}$ given in Proposition 6.26 satisfies the below inequalities.

(a) We have $B_{f,\mathfrak{s}} \leq 4v(2)$. In particular, if $\delta(\mathfrak{s}) \geq 4v(2)$, then there exists a valid disc linked to $\mathfrak{s}$, and if $\delta(\mathfrak{s}) > 4v(2)$, then it is guaranteed that there are exactly 2 valid discs linked to $\mathfrak{s}$.
(b) If $\mathfrak{s}$ has a child cluster (resp. a sibling cluster) of odd cardinality, then we have $B_{f,\mathfrak{s}} \geq \frac{2v(2)}{|\mathfrak{s}|-1}$ (resp. $B_{f,\mathfrak{s}} \geq \frac{2v(2)}{2g+1-|\mathfrak{s}|}$). In particular, if $\mathfrak{s}$ is a cardinality-2 or a cardinality-$2g$ cluster, there cannot be a valid disc linked to $\mathfrak{s}$ if $\delta(\mathfrak{s}) < 2v(2)$.
(c) If $\mathfrak{s}$ has both a child and a sibling cluster of odd cardinality, then we have

$$B_{f,\mathfrak{s}} \geq \left(\frac{2}{|\mathfrak{s}| - 1} + \frac{2}{2g + 1 - |\mathfrak{s}|} \right) v(2).$$

(d) Suppose that we are in one of the following settings.
  (i) Assume that $\mathfrak{s}$ has exactly 2 odd-cardinality child clusters $\mathfrak{r}_1$ and $\mathfrak{r}_2$. Let $\tilde{f}(x) \in \bar{K}[x]$ be a polynomial whose set of roots coincides with $\widetilde{\mathcal{R}} := \mathfrak{r}_1 \sqcup \mathfrak{r}_2 \sqcup (\mathcal{R} \smallsetminus \mathfrak{s})$, and let $\widetilde{\mathfrak{s}} = \mathfrak{r}_1 \sqcup \mathfrak{r}_2$.
  (ii) Assume that $\mathfrak{s}$ has exactly 1 odd-cardinality sibling cluster $\mathfrak{r}_2$, and let $\mathfrak{r}_1$ denote the (odd-cardinality) parent cluster of $\mathfrak{s}$. Let $\tilde{f}(x) \in \bar{K}[x]$ be a polynomial whose set of roots coincides with $\widetilde{\mathcal{R}} := (\mathcal{R} \smallsetminus \mathfrak{r}_1) \sqcup \mathfrak{r}_2 \sqcup \mathfrak{s}$, and let $\widetilde{\mathfrak{s}} = \mathfrak{s}$.

  In each case we have $B_{f,\mathfrak{s}} = B_{\tilde{f},\widetilde{\mathfrak{s}}}$.
(e) Suppose that at least one of the following holds:
  (i) each of the child clusters of $\mathfrak{s}$ has even cardinality and depth $\geq 2v(2)$; or
  (ii) the parent and each of the sibling clusters of $\mathfrak{s}$ have even cardinality and depth $\geq 2v(2)$.

Then we have $B_{f,\mathfrak{s}} \leq 2v(2)$. If both (i) and (ii) above hold, then we have $B_{f,\mathfrak{s}} = 0$.

PROOF. As the continuous piecewise-linear functions $\mathfrak{t}_{\pm}^{\mathfrak{s}}$ have positive integer slopes until reaching an output of $2v(2)$ by Proposition 6.19, we must have $b_0(\mathfrak{t}_{\pm}^{\mathfrak{s}}) \leq 2v(2)$. This implies that $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) + b_0(\mathfrak{t}_{-}^{\mathfrak{s}}) \leq 4v(2)$, proving part (a).

Now if $\mathfrak{s}$ has a child cluster (resp. a sibling cluster) of odd cardinality, then we have $\mathfrak{t}_{+}^{\mathfrak{s}}(0) = 0$ (resp. $\mathfrak{t}_{-}^{\mathfrak{s}}(0) = 0$) by Lemma 6.34. By Lemma 6.33, we then have $b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) \geq \frac{2v(2)}{|\mathfrak{s}|-1}$ (resp. $b_0(\mathfrak{t}_{-}^{\mathfrak{s}}) \geq \frac{2v(2)}{2g+1-|\mathfrak{s}|}$). This proves (b) and (c) if we recall that $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) + b_0(\mathfrak{t}_{-}^{\mathfrak{s}})$.

Let us now address (d). Taking into account Remark 6.29, the alternate hypotheses of part (d) correspond respectively to the hypotheses of parts (a) and (c) of Proposition 6.28(a), with $\mathfrak{r}$ being $\mathfrak{r}_1 \sqcup \mathfrak{r}_2$ in case (i), and with $\mathfrak{r}$ being the union of $\mathfrak{s}$ with its even-cardinality sibling clusters in case (ii). Let us first address case (i). By applying Proposition 6.28, we get $\mathfrak{t}_{+}^{\mathfrak{s}}(b) = \min\{b, 2v(2)\}$ for all $b \in [0, +\infty)$. Moreover, we have that $\tilde{\mathfrak{s}} := \mathfrak{r}$ is an even-cardinality cluster of the set $\tilde{\mathcal{R}} := \mathfrak{r} \sqcup (\mathcal{R} \smallsetminus \mathfrak{s})$. Now, since $\mathcal{R} \smallsetminus \mathfrak{s} = \widetilde{\mathcal{R}} \smallsetminus \tilde{\mathfrak{s}}$, we have that the two functions $t_{-}^{\mathfrak{s}}$ and $t_{-}^{\tilde{\mathfrak{s}}}$ coincide; on the other hand, Proposition 6.28(a) can also be applied to $\tilde{\mathfrak{s}} \subseteq \widetilde{\mathcal{R}}$ to get that $\mathfrak{t}_{+}^{\tilde{\mathfrak{s}}}(b) = \min\{b, 2v(2)\}$. We conclude that $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) + b_0(\mathfrak{t}_{-}^{\mathfrak{s}}) = b_0(\tilde{f}_{+}^{\tilde{\mathfrak{s}}}) + b_0(\tilde{f}_{-}^{\tilde{\mathfrak{s}}}) = B_{\tilde{f},\tilde{\mathfrak{s}}}$. The proof for case (ii) is completely analogous.

Let us now address case (i) of (e). Since the child clusters $\mathfrak{c}_1, \ldots, \mathfrak{c}_N$ of $\mathfrak{s}$ have depth $\geq 2v(2)$, we get that $\mathfrak{t}_{+}^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i)) = 2v(2)$ for $1 \leq i \leq N$, due to the fact that each function $\mathfrak{t}_{+}^{\mathfrak{c}_i}$ has positive integer slopes as long as its output is $< 2v(2)$ by Proposition 6.19. Then by Proposition 6.30, we have $\mathfrak{t}_{+}^{\mathfrak{s}}(0) = 2v(2)$ also, which directly implies that $b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) = 0$. In a similar manner, one proves that under assumption (ii), we get $\mathfrak{t}_{-}^{\mathfrak{s}}(0) = 2v(2)$, which directly implies that $b_0(\mathfrak{t}_{-}^{\mathfrak{s}}) = 0$. Now in general, as was observed in the proof of part (a), we have $b_0(\mathfrak{t}_{\pm}^{\mathfrak{s}}) \leq 2v(2)$, and thus the claims of part (e) follow from the formula $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_{+}^{\mathfrak{s}}) + b_0(\mathfrak{t}_{-}^{\mathfrak{s}})$. □

5. Sufficiently odd part-square decompositions

The motivation for this subsection was presented at the end of §6.3.3.3. For this subsection, we will let $h \in K[z]$ be a nonzero polynomial whose roots (in $\bar{K}$) all have valuation ≤ 0: the cases we care about are those of $h = f_{\pm}^{\mathfrak{s},\alpha}$, where $f_{\pm}^{\mathfrak{s},\alpha}$ are the two functions introduced in §6.3.3.2, whose roots satisfy this property according toRemark 6.22(b). Let us choose a part-square decomposition $h = q^2 + \rho$, and let us consider the function $[0, +\infty) \ni b \mapsto \underline{t}_{q,\rho}(D_{0,b}) = \underline{v}_{\rho}(D_{0,b}) - \underline{v}_h(D_{0,b})$.

Remark 6.36. The assumption on h easily implies that the piecewise-linear function $[0, +\infty) \ni b \mapsto \underline{v}_h(D_{0,b})$ is constant (see Lemma 6.2); as a consequence, the linear function $[0, +\infty) \ni b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ is a non-decreasing piecewise-linear function with decreasing slopes and has the same slopes as $[0, +\infty) \ni b \mapsto \underline{v}_{\rho}(D_{0,b})$.

Definition 6.37. A part-square decomposition $h = q^2 + \rho$ of a polynomial $h \in \bar{K}[z]$ whose roots all have valuation ≤ 0 is said to be *sufficiently odd* if the function $\underline{t}_{q,\rho} : [0, +\infty) \ni b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ satisfies the following.

(a) We have $\underline{t}_{q,\rho}(D_{0,b}) \geq 2v(2)$ for some $b \in [0, +\infty)$; we will denote by $b_0(\underline{t}_{q,\rho})$ the minimal $b \in [0, +\infty)$ with this property, so that $\underline{t}_{q,\rho}(D_{0,b}) < 2v(2)$ for $b \in [0, b_0(\underline{t}_{q,\rho}))$, and $\underline{t}_{q,\rho}(D_{0,b}) \geq 2v(2)$ for $b \in [b_0(\underline{t}_{q,\rho}), +\infty)$.
(b) If $b_0(\underline{t}_{q,\rho}) > 0$, the left derivative of $b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ at $b = b_0(\underline{t}_{q,\rho})$ is odd.

Remark 6.38. We have the following.

(a) Every part-square decomposition in which ρ has no constant term satisfies condition (a) of Definition 6.37, because when ρ has no constant term, the function $b \in [0, +\infty) \mapsto \underline{v}_\rho(D_{0,b})$ cannot have slope 0 (see Lemma 6.2), and hence $b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ is a strictly increasing function.
(b) Totally odd decomposition are sufficiently odd, because for a totally odd decomposition all the slopes of $[0, +\infty) \ni b \mapsto \underline{v}_\rho(D_{0,b})$, and hence all the slopes of $[0, +\infty) \ni b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ over $[0, +\infty)$, are odd. In particular, Proposition 4.20 implies that a sufficiently odd decomposition always exists.

Proposition 6.39. If $h = q^2 + \rho$ is a sufficiently odd part-square decomposition, then it is good at all discs $D_{0,b}$, with $b \in [\max\{0, b_0(\underline{t}_{q,\rho}) - \varepsilon\}, +\infty)$, for $\varepsilon > 0$ small enough.

Proof. For $b \in [b_0(\underline{t}_{q,\rho}), +\infty)$, we have $\underline{t}_{q,\rho}(D_{0,b}) \geq 2v(2)$, which clearly implies that the decomposition is good at $D_{0,b}$. Let us therefore assume that $b_0(\underline{t}_{q,\rho}) > 0$ and focus on the interval $[b_0(\underline{t}_{q,\rho}) - \varepsilon, b_0(\underline{t}_{q,\rho}))$, where $\varepsilon > 0$ has been chosen small enough that the function $b \mapsto \underline{t}_{q,\rho}(D_{0,b})$ is $< 2v(2)$ has odd slope, as prescribed by Definition 6.37. Via Lemma 6.2, we deduce that any normalized reduction of $\rho_{0,\beta}$, for $\beta \in \bar{K}^\times$ any element of valuation $b \in [b_0(\underline{t}_{q,\rho}) - \varepsilon, b_0(\underline{t}_{q,\rho}))$, is not a square, hence the decomposition of h is good at $D_{0,b}$ thanks to Proposition 4.18(a). $\square$

Corollary 6.40. The value of $b_0(\underline{t}_{q,\rho})$ is the same for all sufficiently odd decompositions $h = q^2 + \rho$ of the polynomial h. When $b_0(\underline{t}_{q,\rho}) > 0$, the left derivative $\lambda(\underline{t}_{q,\rho})$ of the function $b \in [0, +\infty) \mapsto \underline{t}_{q,\rho}(D_{0,b})$ at $b = b_0(\underline{t}_{q,\rho})$ (which is an odd positive integer) is also independent of the sufficiently odd decomposition.

Proof. This follows immediately from the proposition above, taking into account Remark 4.17. $\square$

Now we recall the the invariants $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ and $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ introduced in Subsection 6.3.3.3 for any even-cardinality-subset $\mathfrak{s} \subseteq \mathcal{R}$ and $\alpha \in \bar{K}$, the knowledge of which is sufficient to determine the sub-interval $J(\mathfrak{s}, \alpha) \subseteq I(\mathfrak{s}, \alpha)$ and the slopes $\lambda_\pm(\mathfrak{s}, \alpha)$ that we have introduced in §6.2 and that play a crucial role in determining which are the valid discs D centered at α as well as the corresponding structure of $(\mathcal{Y}_D)_s$. In the language of this subsection, the quantities $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ and $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ are nothing but $b_0(\underline{t}_{q_\pm,\rho_\pm})$ and $\lambda(\underline{t}_{q_\pm,\rho_\pm})$ for two totally odd part-square decompositions $f_\pm^{\mathfrak{s},\alpha} = q_\pm^2 + \rho_\pm$, where $f_\pm^{\mathfrak{s},\alpha}$ are the two polynomials introduced in Subsection 6.3.3.2: this is clear from Remark 6.22(c). Now, the following proposition immediately follows from Corollary 6.40.

Proposition 6.41. We have $b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha}) = b_0(\underline{t}_{q_\pm,\rho_\pm})$ and $\lambda(\mathfrak{t}_\pm^{\mathfrak{s},\alpha}) = \lambda(\underline{t}_{q_\pm,\rho_\pm})$ for any two *sufficiently odd* (and not necessarily totally odd) part-square decompositions $f_\pm^{\mathfrak{s},\alpha} = q_\pm^2 + \rho_\pm$.

We conclude by observing that, for a valid disc D, the equation $\mathcal{Y}_D \to \mathcal{X}_D$ can also be written down from the knowledge of sufficiently odd decompositions only.

Proposition 6.42. Let α and $\mathfrak{s}$ be as in the assumptions of Theorem 6.18; assume that we have $J(\mathfrak{s}, \alpha) \neq \varnothing$, and let $D := D_{\alpha, b_+(\mathfrak{s},\alpha)}$ or $D := D_{\alpha, b_-(\mathfrak{s},\alpha)}$ be a valid disc provided by the theorem. Then the part-square decomposition of f we obtain from any two chosen sufficiently odd decompositions of $f_+^{\mathfrak{s},\alpha}$ and $f_-^{\mathfrak{s},\alpha}$ (see §6.3.3.2) is good at the disc D.

Proof. Let $f_\pm^{\mathfrak{s},\alpha} = q_\pm^2 + \rho_\pm$ be any sufficiently odd decompositions for $f_+^{\mathfrak{s},\alpha}$ and $f_-^{\mathfrak{s},\alpha}$, and let $f^{\mathfrak{s}} = (q^{\mathfrak{s}})^2 + \rho^{\mathfrak{s}}$, $f^{\mathcal{R} \smallsetminus \mathfrak{s}} = (q^{\mathcal{R} \smallsetminus \mathfrak{s}})^2 + \rho^{\mathcal{R} \smallsetminus \mathfrak{s}}$ and $f = q^2 + \rho$ be the decompositions they induce for $f^{\mathfrak{s}}$, $f^{\mathcal{R} \smallsetminus \mathfrak{s}}$ and f (see §6.3.3.2). By Definition 6.37 and Proposition 6.41, we have $\underline{t}_{q_\pm, \rho_\pm}(D_{0,b}) \geq 2v(2)$ when $b \geq b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$. Now, Remark 6.22(a) ensures that the corresponding decompositions $f^{\mathfrak{s}} = (q^{\mathfrak{s}})^2 + \rho^{\mathfrak{s}}$ and $f^{\mathcal{R} \smallsetminus \mathfrak{s}} = (q^{\mathcal{R} \smallsetminus \mathfrak{s}})^2 + \rho^{\mathcal{R} \smallsetminus \mathfrak{s}}$ satisfy $\underline{t}_{q^{\mathfrak{s}}, \rho^{\mathfrak{s}}}(D_{\alpha,b}) \geq 2v(2)$ for $b \leq b_+(\mathfrak{s}, \alpha)$, and $\underline{t}_{q^{\mathcal{R} \smallsetminus \mathfrak{s}}, \rho^{\mathcal{R} \smallsetminus \mathfrak{s}}}(D_{\alpha,b}) \geq 2v(2)$ for $b \geq b_-(\mathfrak{s}, \alpha)$, recalling that $b_\pm(\mathfrak{s}, \alpha) = d_\pm(\mathfrak{s}, \alpha) \mp b_0(\mathfrak{t}_\pm^{\mathfrak{s},\alpha})$ by Remark 6.25. Since we have $J(\mathfrak{s}, \alpha) \neq \varnothing$, which is to say that $b_-(\mathfrak{s}, \alpha) \leq b_+(\mathfrak{s}, \alpha)$ (see Remark 6.25), we consequently have $\underline{t}_{q^{\mathfrak{s}}, \rho^{\mathfrak{s}}}(D_{\alpha,b}) \geq 2v(2)$ and $\underline{t}_{q^{\mathcal{R} \smallsetminus \mathfrak{s}}, \rho^{\mathcal{R} \smallsetminus \mathfrak{s}}}(D_{\alpha,b}) \geq 2v(2)$ for $b \in \{b_+(\mathfrak{s}, \alpha), b_-(\mathfrak{s}, \alpha)\}$. By Proposition 4.21(a), we have $\underline{t}_{q,\rho}(D_{\alpha,b}) \geq 2v(2)$ for those values of b, i.e. the part-square decomposition for f is good at D. □

6. An algorithm for finding sufficiently odd part-square decompositions

As previously explained, our main motivation for using sufficiently odd part-square decompositions rather than totally odd ones is that it is generally easier to compute a sufficiently odd decomposition of a polynomial. The following is a general algorithm for finding sufficiently odd part-square decompositions of a nonzero polynomial $h(z) \in K[z]$ whose roots all have valuation ≤ 0; more precisely, starting with some part-square decomposition $h = q^2 + \rho$ with $v(\rho) \geq v(h)$, this algorithm (when it terminates) transforms it into a sufficiently odd part-square decomposition of h, essentially by modifying q by adding square roots of even-degree terms of ρ. (We note that this algorithm is very similar to the procedure given by [13, Proposition 2.2.1], which has a similar aim, although the latter involves adding the square roots of all even-degree terms simultaneously at each step and in that way is dissimilar to our method.)

Algorithm 6.43. Let $h(z) \in K[z]$ be a nonzero polynomial whose roots all have valuation ≤ 0, and let $h = q^2 + \rho$ be a part-square decomposition of h satisfying $v(\rho) \geq v(h)$ (for instance, we may choose the trivial decomposition $h = 0^2 + \rho$). In the steps below, we will change the polynomial $q(z)$ without changing $h(z)$, modifying $\rho(z)$ accordingly so that $h = q^2 + \rho$ is always a part-square decomposition of h. At each stage, we write R_i for the ith coefficient of ρ.

(1) Choose some ordering $n_0, n_1, \ldots, n_{\lfloor \frac{1}{2} \deg(h) \rfloor}$ of the set of natural numbers $\{0, 1, \ldots, \lfloor \frac{1}{2} \deg(h) \rfloor\}$. In practice, this algorithm produces cleaner and more efficient results when we let $n_0 = 0$ and use the following ordering for the natural numbers in $\{1, \ldots, \lfloor \frac{1}{2} \deg(h) \rfloor\}$. Each natural number can be written uniquely as $s2^j$ for a positive odd integer s and an integer $j \geq 0$. Then a natural number $s2^j$ comes before another natural number $s'2^{j'}$ in our ordering if and only if either we have $s < s'$

or we have $s = s'$ and $j' > j$; in other words, we order these numbers first according to their maximal odd factors and then in *descending* order of their maximal 2-power factors.

Now for $0 \leq i \leq \lfloor \frac{1}{2} \deg(h) \rfloor$, perform the following two steps:

(i) Replace $q(z)$ with $q(z) + \sqrt{R_{2n_i}} z^{n_i}$ and modify $\rho(z)$ accordingly.

(ii) Check whether the decomposition $h = q^2 + \rho$ is a sufficiently odd part-square decomposition of f, and if it is, terminate the algorithm.

(2) Repeat Step (1).

The next results show that the above algorithm terminates after a finite number of steps under certain hypotheses.

Lemma 6.44. Assume the set-up and notation in Algorithm 6.43. Suppose that we have completed Step (1) of Algorithm 6.43 a total of N times for some $N \geq 0$, ignoring Step (1)(ii) (in other words, performing Step (1)(i) for a_i ranging through all natural numbers in $\{1, \ldots, \lfloor \frac{1}{2} \deg(h) \rfloor\}$) on the Nth time. For positive integers $j \leq \frac{1}{2} \deg(h)$, we have $v(R_{2j}) - v(h) \geq 2v(2)(1 - 2^{-N})$. Moreover, if $N \geq 1$ and the suggested ordering of the n_i's has been used in each rendition of Step 1, then we have $R_0 = R_2 = 0$.

Proof. We prove this claim inductively, starting with the fact that it obviously holds for $N = 0$ as in this case we have $2v(2)(1 - 2^{-N}) = 0$ and we have $v(R_0) = v(h)$ since the roots all have valuation ≤ 0. Now assume that the claim holds for some $N \geq 0$ and consider how our part-square decomposition changes as we perform Step (1) for the $(N+1)$th time. Since all even-power terms of ρ have valuation at least $2v(2)(1 - 2^{-N}) + v(h)$, it is easy to see from the instructions of Step (1)(i) that the terms we are adding to $q(z)$ all have valuations at least $v(2)(1 - 2^{-N}) + \frac{1}{2}v(h)$. Meanwhile, each power-$2j$ term of $\rho(z)$ is eliminated at the i_0th rendition of Step (1)(i) where $n_{i_0} = j$ and may only reappear during a later rendition of Step (1)(i) (the ith rendition for some $i > i_0$) as 2 times the product of two terms of $q(z)$, one of which has been newly added: these are the x^a- and x^b-terms in $q(z)$ for some $a, b \geq 0$ with $a \neq b$ and $a + b = 2j$. Note that if the suggested ordering of the n_i's is followed, then this later n_i cannot equal $2j$ and so we even have $a, b \geq 1$ in this case. Therefore the coefficient of this new power-$2j$ term of $\rho(z)$ has valuation at least

$$\left[v(2) + \frac{1}{2}v(h)\right] + \left[v(2)(1 - 2^{-N}) + \frac{1}{2}v(h)\right] = 2v(2)(1 - 2^{-(N+1)}) + v(h). \tag{34}$$

This proves the claim for $N+1$. Meanwhile, if the suggested ordering of the n_i's has been followed, the numbers a and b defined above, being distinct positive integers whose sum is an even number, must satisfy $a + b \geq 4$. It follows that in this case, $\rho(z)$ has no constant or quadratic term (in other words, $R_0 = R_2 = 0$) after any number $N \geq 1$ of repetitions of Step (1). □

Proposition 6.45. Assume that the coefficient of the x^s-term of $h(z)$ has valuation equal to $v(h)$ for some odd integer $s \geq 1$. Then Algorithm 6.43 terminates after repeating Step (1) at most $\max\{1, \lfloor \log_2(s) \rfloor - 1\}$ times if the suggested ordering of the n_i's is followed.

Proof. First of all, since the square of any polynomial in $g(z) \in K[z]$ has the property that its odd-degree coefficients have valuation at least $v(2) + v(g)$, at any point while running the algorithm, it is clear that the z^s-coefficient R_s in $\rho(z) = h(z) - q^2(z)$ still has valuation equal to $v(h)$.

By Remark 6.38(a) and the last statement of Lemma 6.44, since at any point after Step (1) has been performed the first time the polynomial $\rho(z)$ has no constant term, the criterion in Definition 6.37(a) is satisfied at any point after the first rendition of Step (1). Assume that Step (1) has just been performed for the ($N := \max\{1, \lfloor \log_2(s) \rfloor - 1\}$)th time in the course of running Algorithm 6.43 and that the suggested ordering has always been followed; we shall show that the decomposition $h = q^2 + \rho$ that we have constructed by this point is sufficiently odd, which implies the statement of the proposition. Note that the definition of the integer N directly implies the inequality $N \geq \log_2(s+1) - 2$. It now follows directly from Lemma 6.44 that for integers j such that $1 \leq j \leq \lfloor \frac{1}{2} \deg(h) \rfloor$, we have

$$v(R_{2j}) - v(h) \geq 2v(2)(1 - 2^{-N}) \geq 2v(2)(1 - 2^{-\log_2(s+1)+2}) = \big(2 - \frac{8}{s+1}\big)v(2). \tag{35}$$

As in Definition 6.37(a), let $b_0 := b_0(\underline{t}_{q,\rho}) \in [0, +\infty)$ be the (unique) minimal non-negative rational number satisfying $\underline{t}_{q,\rho}(b) \geq 2v(2)$, and choose an element $\beta_0 \in \bar{K}^\times$ with $v(\beta_0) = b_0$. As we have $N \geq 1$, we also have $R_0 = R_2 = 0$ by Lemma 6.44, so in particular the polynomial $\rho(\beta_0 z)$ does not have a constant or quadratic term whose coefficient has valuation equal to $v(\rho(\beta_0 z))$. We therefore assume that $j \geq 2$ and proceed to show that the coefficient of the power-$2j$ term of $v(\rho(\beta_0 z))$ does not have valuation equal to $v(h(\beta_0 z))$ either; from Definition 6.37 and Lemma 6.2(b), this implies that $h = q^2 + \rho$ is sufficiently odd and the proposition will be proved. Now from equation (35) we have

$$2v(2) - v(R_{2j}) + v(h) \leq \frac{8}{s+1}v(2) \leq \frac{4j}{s+1}v(2) < \frac{4j}{s}v(2). \tag{36}$$

Our assumption that the roots of h each have valuation ≤ 0 implies that $\underline{v}_h(b)$ is constant for $b \in [0, +\infty)$ and equal to $v(h)$, so in fact we have $v(\rho(\beta_0 z)) = \underline{v}_\rho(b_0) = 2v(2) + \underline{v}_h(b_0) = 2v(2) + v(h)$. Our hypothesis on the z^s-coefficient now tells us that $sb_0 = v(\beta_0^s R_s) - v(h) \geq v(\rho(\beta_0 z)) - v(h) = 2v(2)$ and therefore we have $b_0 \geq \frac{2}{s}v(2)$. Now, using (36), we get

$$v(\beta_0^{2j} R_{2j}) = v(R_{2j}) + 2jb_0 > 2v(2) + v(h) - \frac{4j}{s}v(2) + \frac{4j}{s}v(2) = 2v(2) + v(h), \tag{37}$$

which proves the desired statement.

□

Remark 6.46. It is not difficult to show that the algorithm terminates under the first hypothesis of Lemma 6.44 (but not necessarily within $\max\{1, \lfloor \log_2(s) \rfloor - 1\}$ renditions of Step 1) even when the suggested ordering of the n_i's is not followed, through a proof similar to the above one but which does not rely on the conclusion of the last statement of Lemma 6.44, namely that $R_2 = 0$. Similarly, it is evident from the above proof that the algorithm terminates under the weaker condition that the coefficient of the x^s-term has valuation $< v(h) + v(2)$. However, in this slightly more general situation it is much messier to write down a bound for the number of times Step (1) must be performed.

Remark 6.47. One may apply Algorithm 6.43 to get sufficiently odd decomposition for the polynomials $h = f_\pm^{\mathfrak{s},\alpha}$ we have introduced in §6.3.3.2, where $\mathfrak{s} \subseteq \mathcal{R}$ is an even-cardinality subset and $\alpha \in \bar{K}$, as the roots of these polynomials all have valuation ≤ 0

by Remark 6.22(b). However, we note that Proposition 6.45 cannot be applied in general to guarantee that the algorithm terminates, because the condition that there exists an odd integer s such that the coefficient of the power-s term is a unit does not necessarily hold. In fact, it is not difficult to see that, when $\mathfrak{s}$ is a cluster and $\alpha \in D_{\mathfrak{s},d_+(\mathfrak{s})}$, this condition holds for $f_-^{\mathfrak{s},\alpha}$ (resp. $f_+^{\mathfrak{s},\alpha}$) if and only if the cardinality of the child cluster of $\mathfrak{s}$ which contains α (resp. the parent cluster of $\mathfrak{s}$) is an odd integer s. We suspect that the algorithm still terminates under weaker conditions.

7. Computations of sufficiently odd part-square decompositions for low degree

In this subsection we use Algorithm 6.43 to compute general formulas for sufficiently odd part-square decompositions in the cases that the polynomial $h(z) \in K[z]$ has odd degree at most 7 (noting that by construction, in §6.3.3.2 the polynomials $f_{\pm}^{\mathfrak{s},\alpha}$ for which we want sufficiently odd decompositions have odd degree provided that $\alpha \in \mathfrak{s}$), with an additional hypothesis in the case of degree 7 that the roots are all units. Some of the formulas we obtain will be used for the computations in §9.

In the cases of degree 1, 3, and 5, in fact the formulas computed below give us totally odd decompositions, whereas in degree 7, Algorithm 6.43 terminates and gives us formulas for a sufficiently (but not totally) odd decomposition.

7.1. Polynomials of degree 1. It is immediate to see that given a linear polynomial $h(z) \in \bar{K}[z]$, letting $q(z)$ be a square root of $h(0)$ and $\rho(z) = h(z) - q^2(z) = h(z) - h(0)$, the decomposition $h = q^2 + \rho$ is sufficiently (and even totally) odd.

7.2. Polynomials of degree 3. Let $h(z) = \sum_{i=0}^{3} H_i x^i \in \bar{K}[z]$ be a cubic polynomial. Then it is clear that Algorithm 6.43 terminates during the first rendition of Step 1 after performing Step 1(i) for $i = 1$ following the suggested ordering (with $s_0 = 0$ and $s_1 = 1$), and we have $q(z) = \sqrt{H_0} + \sqrt{H_2}z$ (for some choices of square roots of H_0 and H_2) and

$$\rho(z) = (H_1 - 2\sqrt{H_2}\sqrt{H_0})z + H_3 z^3. \tag{38}$$

This decomposition $h = q^2 + \rho$ is therefore sufficiently (and even totally) odd. One easily checks that this is the same totally odd decomposition obtained by using the method described in the proof of Proposition 4.20.

7.3. Polynomials of degree 5. Let $h(z) = \sum_{i=0}^{5} H_i z^i \in \bar{K}[z]$ be a quintic polynomial. In this case once again Algorithm 6.43 terminates during the first rendition of Step 1 after performing Step 1(i) for $i = 2$ following the suggested ordering (with $s_0 = 0$, $s_1 = 2$, and $s_2 = 1$); it is straightforward to check that here we have $q(z) = \sqrt{H_0} + \sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}z + \sqrt{H_4}z^2$ (where we have chosen square roots of H_0 and H_4 and then chosen a square root of $H_2 - 2\sqrt{H_4}\sqrt{H_0}$) and

$$\rho(z) = \left(H_1 - 2\sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}\sqrt{H_0}\right)z + \left(H_3 - 2\sqrt{H_4}\sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}\right)z^3 + H_5 z^5. \tag{39}$$

We have therefore again found a sufficiently (and even totally) odd decomposition $h = q^2 + \rho$. It is in fact not too difficult to show that (similarly to the $g = 1$ case) we obtain this same totally odd decomposition by using the method described in the proof of Proposition 4.20.

7.4. Polynomials of degree 7 with unit roots. Let $h(z) = \sum_{i=0}^{7} H_i z^i \in \bar{K}[z]$ be a septic polynomial whose roots all have valuation 0 (so that in particular we have $v(H_7) = v(h)$). Now according to Proposition 6.45, Step 1 of Algorithm 6.43 needs to performed only $\max\{1, \lfloor \log_2(7) - 1 \rfloor\} = 1$ time. In our only rendition of Step 1, following the suggested ordering (with $s_0 = 0$, $s_1 = 2$, and $s_2 = 1$, and $s_3 = 3$), we at most need to perform Step 1(i) for $i = 0, 1, 2, 3$ in order to obtain a sufficiently odd decomposition $h = q^2 + \rho$; after doing this, it is straightforward to check that we have

$$q(z) = \sqrt{H_0} + \sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}z + \sqrt{H_4}z^2 + \sqrt{H_6}z^3 \tag{40}$$

(where we have first chosen square roots of H_0, H_4, and H_6 and then chosen a square root of $H_2 - 2\sqrt{H_4}\sqrt{H_0}$) and

(41)

$$\begin{aligned}\rho(z) = &\left(H_1 - 2\sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}\sqrt{H_0}\right)z + \left(H_3 - 2\sqrt{H_4}\sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}} - 2\sqrt{H_6}\sqrt{H_0}\right)z^3 \\ &- 2\sqrt{H_6}\sqrt{H_2 - 2\sqrt{H_4}\sqrt{H_0}}\ z^4 + \left(H_5 - 2\sqrt{H_6}\sqrt{H_4}\right)z^5 + H_7 z^7.\end{aligned}$$

CHAPTER 7

Finding centers of valid discs in the $p = 2$ setting

This section will deal with the problem of determining a center for each valid disc D in the $p = 2$ setting. When $\mathfrak{s} := D \cap \mathcal{R} \neq \varnothing$, the problem is easily solved, since a center of D can be chosen to be any root in $\mathfrak{s} \subseteq \mathcal{R}$. When $\mathfrak{s} = \varnothing$, we will show in §7.2 that D necessarily contain a root of a certain polynomial $F(T) \in K[T]$ that is introduced in §7.1.

1. Defining the polynomial F

Given the hyperelliptic curve $y^2 = f(x)$, with $f(x) \in K[x]$ of odd degree $2g + 1$, Proposition 4.20 allows us to produce (for instance by using the procedure explained in the proof) a totally odd decomposition of the translated polynomial $f_{T,1}(z) := f(z + T)$, in which T remains generic rather than being assigned to be particular center $\alpha \in \bar{K}$. Such a decomposition will have the form

$$f_{T,1} = q_{T,1}^2 + \rho_{T,1},$$

with

$$q_{T,1}(z) = Q_0(T) + Q_1(T)z + \ldots + Q_g(T)z^g \qquad \text{and}$$
$$\rho_{T,1}(z) = R_1(T)z + R_3(T)z^3 + \ldots + R_{2g+1}(T)z^{2g+1},$$

where $Q_i(T)$ and $R_i(T)$ are elements of $\overline{K(T)}$, i.e. algebraic functions of the variable T.

Proposition 7.1. The algebraic functions $Q_i(T)$ and $R_i(T)$ are integral over $K[T]$.

Proof. The proposition is a reflection of the general fact that a totally odd decomposition $h = q^2 + \rho$ of any polynomial h is always good (see Corollary 4.19) and hence, in particular, satisfies $v(\rho) \geq v(h)$ (by Remark 4.16). We now give an explicit proof adapted to the specific setting in which we are working. All we have to show is that, for every valuation subring $\mathcal{O}$ of $\overline{K(T)}$ such that $K[T] \subseteq \mathcal{O}$, the polynomial $q_T(z) \in \overline{K(T)}[z]$ has coefficients in $\mathcal{O}$. Let w be the valuation of $\overline{K(T)}$ whose ring of integers is $\mathcal{O}$; for any polynomial $h(z) \in \overline{K(T)}[z]$, let us also denote the Gauss valuation of h by $w(h)$, i.e. $w(h)$ is the minimum of the valuations of the cofficients of h, and let k_w denote the residue field of w. Suppose by way of contradiction that we have $w(q_{T,1}) < 0$; then, since $f_{T,1} = q_{T,1}^2 + \rho_{T,1}$ and $w(f_{T,1}) \geq 0$, we necessarily have that $w(\rho_{T,1}) = w(q_{T,1}^2) < w(f_{T,1})$. Let $\gamma \in \overline{K(T)}$ be any element such that γ^2 has valuation equal to $w(\rho_{T,1}) = w(q_{T,1}^2)$; now if we multiply the equation $f_{T,1} = q_{T,1}^2 + \rho_{T,1}$ by γ^{-2} and we reduce, we obtain the equation $\overline{\gamma^{-2}\rho_{T,1}} = -(\overline{\gamma^{-1}q_{T,1}})^2$ in $k_w[z]$. But this is impossible, since the left-hand side is a nonzero polynomial with coefficients in $k(w)$ whose monomials all have odd degree, while the right-hand side is the square of a nonzero polynomial with coefficients in $k(w)$, and hence it will contain nonzero monomials of even degree. ∎

Definition 7.2. Let $L \subset \overline{K(T)}$ be the smallest Galois extension of $K(T)$ to which $R_1(T)$ belongs. We define $F(T) \in K[T]$ to be the norm of $R_1(T)$ with respect to the extension $L/K(T)$.

Remark 7.3. Note that we can be sure that the norm $F(T)$ of $R_1(T)$ is actually a polynomial in the variable T (and not just a rational function) because of the integrality result given by Proposition 7.1.

Remark 7.4. In the cases of $g \in \{1, 2\}$, assuming, for simplicity, that f is monic, we may easily compute $F(T)$ as the norm of $R_1(T)$ using the formulas found in §6.7.7.2,6.7.7.3. For $0 \leq i \leq 2g+1$, let $P_i(T) \in K[T]$ be the z^i-coefficient of $f(z+T) \in K[T][z]$. Then for $g = 1$, we have the formula

$$F = P_1^2 - 4P_2P_0, \tag{42}$$

and for $g = 2$, we have the formula

$$F = (P_1^2 - 4P_2P_0)^2 - 64P_4P_0^3. \tag{43}$$

2. Using the polynomial F to find centers

We will now establish some properties of F; in particular, we will show that each root of F is the center of a valid disc, and that all valid discs D such that $D \cap \mathcal{R} = \varnothing$ contain a root of F.

Proposition 7.5. Let α be a root of F in $\bar{K}$. Then we have the following.

(a) There exists a part-square decomposition $f = q^2 + \rho$ which is totally odd at the center α such that $\rho_{\alpha,1}$ has no linear term.
(b) The element α is not a root of f (i.e. $\alpha \notin \mathcal{R}$), and there exists a valid disc D containing α and such that $\ell(\mathcal{X}_D, \infty) > 0$.
(c) If $D = D_{\alpha,b}$ is minimal among the valid discs satisfying the conditions described in (b), then we have $\ell(\mathcal{X}_D, \overline{x_{\alpha,\beta}} = 0) = 0$ (for a choice of $\beta \in \bar{K}^\times$ with $v(\beta) = b$).

Proof. Statement (a) follows from the definition of the polynomial F as the norm of the linear coefficient of $\rho_{T,1}$. More precisely, let $L' \subset \overline{K(T)}$ be a Galois extension of $K(T)$ to which all the coefficients of the polynomials $q_{T,1}(z)$ and $\rho_{T,1}(z)$ belong, and let S' be the integral closure of $K[T]$ in L'. The norm $\mathrm{Nm}_{L'/K(T)}(R_1(T)) \in K[T]$ must be a power F^d of F; hence, the element α is a root of it. Let us also choose an extension $\widetilde{\psi}_\alpha : S' \to \bar{K}$ of the evaluation map $\psi_\alpha : K[T] \to \bar{K}, T \mapsto \alpha$: it is clear that the polynomials $q_{\alpha,1}(z) := \widetilde{\psi}_\alpha(q_{T,1}(z)) \in \bar{K}[z]$ and $\rho_{\alpha,1}(z) := \widetilde{\psi}_\alpha(\rho_{T,1}(z)) \in \bar{K}[z]$ provide a totally odd part-square decomposition for $f_{\alpha,1}$ whose linear coefficient is $\widetilde{\psi}_\alpha(R_1(T)) \in \bar{K}$. On the other hand, since α is a root of F^d, we have $\widetilde{\psi}_\alpha(F^d) = \psi_\alpha(F^d) = 0$, while at the same time, we have the formula $F^d = \prod_\sigma \sigma(R_1(T))$ as σ ranges among the elements of $\mathrm{Gal}_{L'/K(T)}$. We conclude that $\widetilde{\psi}_\alpha(\sigma(R_1(T)) = 0$ for some $\sigma \in \mathrm{Gal}_{L'/K(T)}$; after replacing $\widetilde{\psi}_a$ with $\widetilde{\psi}_\alpha \circ \sigma$, we may assume that $\sigma = 1$, so that $\widetilde{\psi}_\alpha(R_1(T)) = 0$; part (a) is thus proved.

Let us now prove part (b). Using part (a), we have a part-square decomposition $f = q^2 + \rho$ that is totally odd with respect to the center α and such that the linear term of $\rho_{\alpha,1}$ is zero. Suppose that α is a root of f, so that the polynomial $f_{\alpha,1}$ has no constant term. Then, coming from the fact that $x_{\alpha,1}^3 | \rho_{\alpha,1}(x_{\alpha,1})$, it is easy to see that we have $x_{\alpha,1} | q_{\alpha,1}(x_{\alpha,1})$; it immediately follows that we have $x_{\alpha,1}^2 | f_{\alpha,1}(x_{\alpha,1})$, which contradicts the fact that f has no multiple roots. This proves the first claim of part (b).

Now let us study the function $\mathbb{Q} \ni c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$. When $c \to -\infty$, it is constantly zero, since f has odd degree and $f = 0^2 + f$ is consequently a good part-square decomposition at large enough discs, while when $c \to +\infty$, it is constantly $2v(2)$ since α is not a root of f (this was already mentioned in Remark 6.12). As a consequence, there exists $b \in \mathbb{Q}$ such that the output $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$ is $< 2v(2)$ right before $c = b$ and equals $2v(2)$ at $c \geq b$. Let $\mathfrak{s} = D_{\alpha,b} \cap \mathcal{R}$, so that $d_-(\mathfrak{s}, \alpha) < b = b_-(\mathfrak{s}, \alpha) \leq b_+(\mathfrak{s}, \alpha)$ in the language of §6.2. If $\mathfrak{s} \neq \varnothing$, Theorem 6.18 ensures that the disc $D := D_{\alpha,b}$ is valid. If $\mathfrak{s} = \varnothing$, from the fact that $\rho_{\alpha,1}$ has no linear term we deduce that the function $I(\varnothing, \alpha) \to [0, 2v(2)], c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$, which can be computed as $c \mapsto \min\{\underline{t}_{q,\rho}(D_{\alpha,c}), 2v(2)\}$, grows with slopes ≥ 3 until reaching $2v(2)$ at $c = b$ (in other words, it cannot admit slope 1); hence Theorem 6.18 still guarantees that the disc D is valid since $\lambda_-(\varnothing, \alpha) \geq 3$.

It is an immediate consequence of Lemma 6.13 that we have $\ell(\mathcal{X}_D, \overline{x_{\alpha,\beta}} = 0) = 0$ and $\ell(\mathcal{X}_D, \overline{x_{\alpha,\beta}} = \infty) > 0$, for $\beta \in \bar{K}^\times$ an element of valuation b. In particular, the proof of (b) is finished. Moreover, if we take $D' = D_{\alpha,b'}$ to be any other valid disc centered at α such that $\ell(\mathcal{X}_{D'}, \infty) > 0$, by Lemma 6.13 we must have that the output of $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$ is $< 2v(2)$ for c slightly smaller than b' and $= 2v(2)$ at $c = b'$. But this implies, by the construction of b, that we have $b' \leq b$, i.e. $D' \supseteq D$. This shows that the valid disc D that we found above is the minimal one satisfying the conditions given by part (b) and thus completes the proof of part (c). □

The following theorem provides a statement that is somehow converse to the one of Proposition 7.5 above. Together with that proposition, it is essentially a generalization of [10, Theorem 5.1] (which treats only the geometrically equidistant case), and the underlying strategy of its proof is inspired by that of Lehr and Matignon.

THEOREM 7.6. Suppose that $D = D_{\alpha,b}$ is a valid disc such that $\ell(\mathcal{X}_D, \infty) > 0$, i.e. such that $(\mathcal{Y}_D)_s$ has only one branch above $\infty \in (\mathcal{X}_D)_s$. Then D contains a root of F.

Remark 7.7. Let $D = D_{\alpha,b}$ be a disc, and let $\mathfrak{s} = D \cap \mathcal{R}$, so that we have $b \in (d_-(\mathfrak{s}, \alpha), d_+(\mathfrak{s}, \alpha)]$. Then the disc D satisfies the hypothesis in the above theorem if and only if $b = b_-(\mathfrak{s}, \alpha) \leq b_+(\mathfrak{s}, \alpha)$: this is an easy consequence of Theorem 6.18 together with Proposition 6.17(a). Therefore, a valid disc D does not satisfy the hypothesis in the above theorem if and only if $b = b_+(\mathfrak{s}, \alpha) > b_-(\mathfrak{s}, \alpha)$.

In particular, D always satisfies the hypothesis of the theorem if it is linked to no cluster (i.e., $\mathfrak{s} = \varnothing$), or if it is linked to a unique cluster and is the only disc linked to it (i.e. $\mathfrak{s} \neq \varnothing$ and $d_-(\mathfrak{s}) < b_-(\mathfrak{s}) = b = b_+(\mathfrak{s}) < d_+(\mathfrak{s})$).

Corollary 7.8. Each root of F lies in a valid disc. Conversely, suppose that a valid disc D satisfies one of the two assumptions below:

(a) the disc D is linked to no cluster; or
(b) the disc D is linked to a unique cluster $\mathfrak{s}$, is the unique valid disc linked to $\mathfrak{s}$, and is minimal among valid discs.

Then the disc D contains a root α of F. Moreover, for any such α and for all $a \in D \cap \mathcal{R}$, we have $v(a - \alpha) = b$, where b is the depth of the disc D.

Proof. All roots of F belong to valid discs by Proposition 7.5. Conversely, suppose that D is a valid disc. If D is linked to no cluster, or if it is linked to only one cluster and it is the unique valid disc linked to it, then by Remark 7.7 it satisfies the assumption of Theorem 7.6 (i.e., $\ell(\mathcal{X}_D, \infty) > 0$) and consequently contains a root α of F. Now suppose that D satisfies condition (b), so that we have $d_-(\mathfrak{s}) < b_-(\mathfrak{s}) = b = b_+(\mathfrak{s}) < d_+(\mathfrak{s})$ (see the results in §6.2), and so that moreover, if $P \neq \infty$ is the point of $(\mathcal{X}_D)_s$ to which $\mathfrak{s}$ reduces, we have $\ell(\mathcal{X}_D, P) = 1 + \lambda_+(\mathfrak{s}) > 0$ by Proposition 6.17. But Proposition 7.5(c) ensures that, if $P' \neq \infty$ is the point of $(\mathcal{X}_D)_s$ to which $x = \alpha$ reduces, we have $\ell(\mathcal{X}_D, P') = 0$. Hence, we have $P \neq P'$, which means that $v(a - \alpha) = b$ for all $a \in \mathfrak{s}$. $\square$

Let us now address the proof of Theorem 7.6. We begin with the following lemma.

Lemma 7.9. Let $D = D_{\alpha,b}$ be a valid disc satifying the hypothesis in Theorem 7.6. Then, for small enough $\varepsilon > 0$, we have the following: for all α' such that $b' := v(\alpha-\alpha') \in [b-\varepsilon, b)$, and for any pair of part-square decompositions $f = q^2 + \rho$ and $f = (q')^2 + \rho'$ which are good at the disc $D' := D_{\alpha,b'} = D_{\alpha',b'}$, we have the comparison $v(R_1) > v(R_1')$ between the respective linear coefficients $R_1, R_1' \in \bar{K}$ of $\rho_{\alpha,1}$ and $\rho'_{\alpha',1}$.

Proof. Let $\mathfrak{s}$ and α be as in Remark 7.7, so that we have $d_-(\mathfrak{s}, \alpha) < b = b_-(\mathfrak{s}, \alpha) \leq b_+(\mathfrak{s}, \alpha)$. In particular, by the results in §6.2 we have $\mathfrak{t}^{\mathcal{R}}(D') < 2v(2)$ and that $(\mathcal{Y}_{D'})_s \to (\mathcal{X}_{D'})_s$ is inseparable for all discs $D' := D_{\alpha,b'}$ with $b' \in [b - \varepsilon, b)$, for $\varepsilon > 0$ small enough. After possibly shrinking ε, we furthermore obtain that, for all such b', we claim that

$$\{0\} \subseteq \mathrm{Ctr}(\mathcal{X}_{D'}, \mathcal{X}^{(\mathrm{rst})}) \subseteq \{0, \}. \tag{44}$$

Indeed, the set $\mathrm{Ctr}(\mathcal{X}_{D'}, \mathcal{X}^{(\mathrm{rst})})$ is just the finite union $\bigcup_{\tilde{\mathcal{X}}} \mathrm{Ctr}(\mathcal{X}_{D'}, \tilde{\mathcal{X}})$, where $\tilde{\mathcal{X}}$ varies among the smooth models of the line dominated by $\mathcal{X}^{(\mathrm{rst})}$. For any such $\tilde{\mathcal{X}}$, if we let $\tilde{D}$ be the disc such that $\tilde{\mathcal{X}} = \mathcal{X}_{\tilde{D}}$, we have three possibilities:

(1) $\tilde{D} \subseteq D$; in this case, since $D \subsetneq D'$, we have $\tilde{D} \subsetneq D'$, so that $\mathrm{Ctr}(\mathcal{X}_{D'}, \tilde{\mathcal{X}}) = \{0\}$; note that this case does actually occur at least once, for $\tilde{D} = D$;
(2) $\tilde{D} \supsetneq D$; in this case, since D' is only slightly larger than D, we may also assume that $\tilde{D} \supsetneq D'$, so that $\mathrm{Ctr}(\mathcal{X}_{D'}, \tilde{\mathcal{X}}) = \{\infty\}$; and
(3) $\tilde{D} \cap D = \varnothing$; in this case, since D' is only slightly larger than D, we may also assume that $\tilde{D} \cap D' = \varnothing$, so that $\mathrm{Ctr}(\mathcal{X}_{D'}, \tilde{\mathcal{X}}) = \{\infty\}$.

By Corollary 3.20, the inclusions in (44) imply that, for $b' \in [b - \varepsilon, b)$, the special fiber $(\mathcal{Y}_{D'})_s$ must be singular above $x_{\alpha,b'} = 0$ but non-singular away from $x_{\alpha,b'} = 0$ and $x_{\alpha,b'} = \infty$. Let us now pick an element $\alpha' \in \bar{K}$ such that $v(\alpha' - \alpha) = b'$ and choose an element $\beta' \in \bar{K}^\times$ of valuation b'. We have $D' = D_{\alpha,b'} = D_{\alpha',b'}$ and that the special fiber $(\mathcal{Y}_{D'})_s$ is non-singular above $x_{\alpha',\beta'} = 0$ (since $x_{\alpha',\beta'} = 0$ corresponds to some point whose $x_{\alpha,\beta'}$-coordinate is neither 0 nor ∞).

Now let us choose two part-square decompositions $f = q^2 + \rho$ and $f = (q')^2 + \rho'$ that are good at the disc D': our aim will be to show the comparison $v(R_1) < (R_1')$ between the linear terms of $\rho_{\alpha,1}$ and $\rho'_{\alpha',1}$ under the assumption that the valuation $b' := v(\alpha' - \alpha)$ satisfies $b' \in [b - \varepsilon, b)$. We will actually show this inequality in three steps, by proving the below for the disc $D' := D_{\alpha,b'}$:

(a) $v(\beta' R_1) > \underline{v}_\rho(D')$;
(b) $v(\beta' R_1') = \underline{v}_{\rho'}(D')$;

(c) $\underline{v}_\rho(D') = \underline{v}_{\rho'}(D')$.

To prove (a), we observe that the inseparable curve $(\mathcal{Y}_{D'})_s$ has the equation

$$y^2 = \overline{\gamma^{-1}\rho_0(x_{\alpha,b'})},$$

where ρ_0 is a normalized reduction of $\rho_{\alpha,\beta'}$ (see §4.6). In light of this, since $\beta' R_1$ is the linear term of $\rho_{\alpha,\beta'}$, (a) simply expresses the fact that $(\mathcal{Y}_{D'})_s$ is singular above $x_{\alpha,b'} = 0$. In a completely analogous way, the equation in (b) expresses the fact that $(\mathcal{Y}_{D'})_s$ is not singular above $x_{\alpha',b'} = 0$. Finally, (c) follows from Remark 4.17. □

Lemma 7.10. Suppose that $h \in K[z]$ is a nonzero polynomial and $D := D_{\alpha,b} \subseteq \bar{K}$ is a disc not containing any of the roots of h in $\bar{K}$. Then we have $v(h(z_0)) = v(h(z_1))$ for all $z_0, z_1 \in D$.

PROOF. Let $a_1, \ldots, a_r$ be the roots of h in $\bar{K}$, so that we can write $h(z) = c\prod_{i=1}^r (z - a_i)$ for some $c \in K^\times$. For each i, since z_0 and z_1 are points of D, while s_i is not, we have $v(z_0 - a_i) = v(z_1 - a_i)$, from which it clearly follows that $v(h(z_0)) = v(h(z_1))$. □

PROOF OF THEOREM 7.6. Let S be the minimal finite Galois extension of $K[T]$ to which $R_1(T)$ belongs, so that $F(T) = \mathrm{Nm}_{S/K[T]}(R_1) = \prod_\sigma \sigma(R_1)$, with the product taken over all $\sigma \in \mathrm{Gal}(S/K[T])$. Now let α' be any point of the annulus $D_\varepsilon \smallsetminus D$, where $D_\varepsilon = D_{\alpha,b-\varepsilon}$ for some $\varepsilon > 0$ chosen small enough so that the conclusion of Lemma 7.9 holds. Let us consider the evaluation maps $\psi_\alpha, \psi_{\alpha'} : K[T] \to \bar{K}$ corresponding to α and α'; for each of them, we make the choice of an extension $\widetilde{\psi}_\alpha, \widetilde{\psi}_{\alpha'} : S \to \bar{K}$; the other possible extensions can be obtained by precomposing with appropriate automorphisms $\sigma \in \mathrm{Gal}(S/K[T])$.

We clearly have that $\widetilde{\psi}_\alpha(R_1)$ and $\widetilde{\psi}_{\alpha'}(R_1)$ are the linear terms of $\rho_{\alpha,1}$ and $\rho'_{\alpha',1}$ for two part-square decompositions $f = q^2 + \rho$ and $f = (q')^2 + \rho'$ which are totally odd with respect to the centers α and α' respectively; in particular, both decompositions are good at any disc containing both α and α' (see Remark 6.5); hence, Lemma 7.9 ensures that $v(\widetilde{\psi_\alpha(R_1)}) > v(\widetilde{\psi_{\alpha'}}(R_1))$. Since this holds for any choices of extensions $\widetilde{\psi}_\alpha, \widetilde{\psi}_{\alpha'}$, we deduce that $v(\widetilde{\psi_\alpha}(\prod_\sigma \sigma(R_1))) > v(\widetilde{\psi_{\alpha'}}(\prod_\sigma \sigma(R_1)))$, which is to say that $v(\psi_\alpha(F)) > v(\psi_{\alpha'}(F))$, which in turn is nothing but the comparison $v(F(\alpha)) > v(F(\alpha'))$.

Now suppose by way of contradiction that D does not contain any root of F. One can clearly find a disc D', with $D \subsetneq D' \subseteq D_\varepsilon$, such that also D' does not contain any root of F. Now, for $\alpha' \in D' \smallsetminus D$, the argument above implies that $v(F(\alpha)) > v(F(\alpha'))$, but, in light of Lemma 7.10, this contradicts the assumption that D' does not contain any root of F. □

CHAPTER 8

The geometry of the special fiber

Our purpose in this section is to use the framework we developed in §6 to glean information about the components of the special fiber of the relatively stable model of the hyperelliptic curve Y, based on knowledge of the relationship between its valid discs and the cluster picture associated to the defining polynomial. In particular, in §8.1 we will compute the toric rank of the special fiber of $\mathcal{Y}^{\mathrm{rst}}$, while in §8.2 we will discuss the abelian rank of its irreducible components.

1. The toric rank

In this subsection, after introducing the notions of a *viable* cluster and an *übereven* cluster, we will prove the following theorem that allows to compute the toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$, and hence, by Proposition 2.6, the toric rank of the special fiber of any semistable model of Y defined over any extension of R.

THEOREM 8.1. The toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$ equals the number of non-übereven viable clusters.

Let us begin by defining viable clusters.

Definition 8.2. We say that a cluster $\mathfrak{s}$ is *viable* if the following are satisfied:

(a) $\mathfrak{s}$ has even cardinality; and
(b) there exist 2 distinct valid discs linked to $\mathfrak{s}$.

Remark 8.3. In the above definition, the results presented in §5.2 that (a) implies (b) (see Remark 5.15), while, in the $p = 2$ case, (b) implies (a) (see Theorem 5.13(a)).

Proposition 8.4. Viable clusters are in one-to-one correspondence with the nodes of $(\mathcal{X}^{(\mathrm{rst})})_s$ over which the cover $(\mathcal{Y}^{\mathrm{rst}})_s \to (\mathcal{X}^{(\mathrm{rst})})_s$ is unramified (i.e., the nodes of $(\mathcal{X}^{(\mathrm{rst})})_s$ that have two distinct inverse images in $(\mathcal{Y}^{\mathrm{rst}})_s$).

PROOF. Suppose that $\mathfrak{s}$ is a viable cluster, and let $D_+ \subsetneq D_-$ be the two valid discs linked to it. It follows from Lemma 6.14 that we have $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$ for all discs D satisfying $D_+ \subsetneq D \subsetneq D_-$; hence, the two lines L_+ and L_- of $(\mathcal{X}^{(\mathrm{rst})})_s$ corresponding to the discs D_+ and D_- intersect at a node $P \in (\mathcal{X}^{(\mathrm{rst})})_s$. We know by Proposition 6.17(a) that $(\mathcal{Y}_{D_\pm})_s$ has two branches above $P \in (\mathcal{X}_{D_\pm})_s$, which implies that $(\mathcal{Y}^{\mathrm{rst}})_s \to (\mathcal{X}^{(\mathrm{rst})})_s$ is unramified above P.

Let us now prove the converse implication. Let P be a node of $(\mathcal{X}^{(\mathrm{rst})})_s$ above which $(\mathcal{Y}^{\mathrm{rst}})_s$ is unramified; let L_- and L_+ the two lines of $(\mathcal{X}^{(\mathrm{rst})})_s$ passing through P; and let $D_\pm$ be the corresponding discs. Since the cover $(\mathcal{Y}^{\mathrm{rst}})_s \to (\mathcal{X}^{(\mathrm{rst})})_s$ is unramified above P, no element of $\mathcal{R} \cup \{\infty\}$ reduces to $P \in (\mathcal{X}^{(\mathrm{rst})})_s$, which is equivalent to saying that no element of $\mathcal{R} \cup \{\infty\}$ reduces to the unique node of $(\mathcal{X}_{\{D_+,D_-\}})_s$. In particular, ∞ lies on one and only one of the two lines L_+ and L_- comprising the special fiber $(\mathcal{X}_{\{D_+,D_-\}})_s$, say $\infty \in L_- \smallsetminus L_+$; this implies, in particular, that we have $D_+ \subsetneq D_-$ by Proposition 4.6. We can now write the decomposition $\mathcal{R} = \mathfrak{s} \sqcup (\mathcal{R} \smallsetminus \mathfrak{s})$, where $\mathfrak{s}$ (resp. $\mathcal{R} \smallsetminus \mathfrak{s}$) consists of the roots whose reductions in $(\mathcal{X}_{\{D_+,D_-\}})_s$ lie on $L_+ \smallsetminus L_-$ (resp. $L_- \smallsetminus L_+$). It is now clear that $\mathfrak{s} = D_+ \cap \mathcal{R}$ is a cluster, to which the two distinct valid discs D_- and D_+ are linked. Since $(\mathcal{Y}^{\mathrm{rst}})_s \to (\mathcal{X}^{(\mathrm{rst})})_s$ is unramified above P, we have that $(\mathcal{Y}_{D_-})_s$ (resp. $(\mathcal{Y}_{D_+})_s$) has two branches above $0 \in (\mathcal{Y}_{D_-})_s$ (resp. $\infty \in (\mathcal{Y}_{D_+})_s$), hence, by Proposition 6.17(b), the cluster $\mathfrak{s}$ must have even cardinality. $\square$

Proposition 8.5. If $\mathfrak{s}$ is a viable cluster corresponding to a node $P \in (\mathcal{X}^{(\mathrm{rst})})_s$ as in Proposition 8.4, then the thickness of each of the 2 nodes lying above P is equal to $(b_+(\mathfrak{s}) - b_-(\mathfrak{s}))/v(\pi)$ (with the notation of §6.2).

Proof. This is straightforward from applying Propositions 4.8 and 3.4(b) to Theorem 6.18. $\square$

We now give the other main definition of this section.

Definition 8.6. An cluster $\mathfrak{s}$ is said to be *übereven* if it is viable and if all of its children clusters are also viable.

Remark 8.7. In the $p \neq 2$ setting, every even-cardinality cluster is viable, and so an übereven cluster is just a cluster whose children are all even; this is the definition of "übereven" used in [7].

Lemma 8.8. Let $\mathfrak{s}$ be a cluster, and let $\mathfrak{c}_1, \ldots, \mathfrak{c}_N$ be its children. If each child $\mathfrak{c}_i$ is viable, then we have $b_+(\mathfrak{s}) = d_+(\mathfrak{s})$ (with the notation of §6.2).

Proof. Since $\mathfrak{c}_i$ is viable, we have $\delta(\mathfrak{c}_i) > B_{f,\mathfrak{s}} \geq b_0(\mathfrak{t}_+^{\mathfrak{c}_i})$, which implies that $\mathfrak{t}_+^{\mathfrak{c}_i}(\delta(\mathfrak{c}_i)) = 2v(2)$. Now Proposition 6.30(a) says that we have $\mathfrak{t}_+^{\mathfrak{s}}(0) = 2v(2)$, which directly implies that $b_+(\mathfrak{s}) = d_+(\mathfrak{s})$. $\square$

Proposition 8.9. The assignement $\mathfrak{s} \mapsto D_{\mathfrak{s},d_+(\mathfrak{s})}$ induces a one-to-one correspondence between the übereven clusters and the valid discs D such that the special fiber $(\mathcal{Y}_D)_s$ is reducible (i.e., $(\mathcal{Y}_D)_s$ consists of 2 rational components).

Proof. Suppose first that D is a valid disc such that $(\mathcal{Y}_D)_s$ is reducible. Then, we know by Theorem 4.32 that the elements of $\mathcal{R} \cup \{\infty\}$ reduce to $N \geq 3$ distinct points of $(\mathcal{X}_D)_s$; we consequently have $D = D_{\mathfrak{s},d_+(\mathfrak{s})}$ for some cluster $\mathfrak{s}$ and that the N points are $\overline{\mathfrak{c}_1}, \ldots, \overline{\mathfrak{c}_{N-1}}, \infty = \overline{\mathcal{R} \smallsetminus \mathfrak{s}}$, where $\mathfrak{c}_1, \ldots, \mathfrak{c}_{N-1}$ are the child clusters of $\mathfrak{s}$ (see Lemma 5.10). From the fact that D is a valid disc we deduce that $d_+(\mathfrak{s}) = d_-(\mathfrak{c}_i)$ is a common endpoint of the intervals $J(\mathfrak{s})$ and $J(\mathfrak{c}_i)$ for all i; in particular, these intervals are non-empty and we have $b_+(\mathfrak{s}) = d_+(\mathfrak{s}) = d_-(\mathfrak{c}_i) = b_-(\mathfrak{c}_i)$. The fact that $(\mathcal{Y}_D)_s$ consists of 2 components means that we have $\ell(\mathcal{X}_D, \overline{\mathfrak{c}_i}) = 0$ for all i and $\ell(\mathcal{X}_D, \infty) = 0$, but, according to Proposition 6.17(c),(d), this implies that $b_-(\mathfrak{c}_i) < b_+(\mathfrak{c}_i)$ for all i as well as $b_-(\mathfrak{s}) < b_+(\mathfrak{s})$. Now, applying Proposition 6.17(b), we also deduce that $\mathfrak{s}$ as well as the $\mathfrak{c}_i$'s must all have even cardinality. We conclude that $\mathfrak{s}$ is a übereven cluster.

Let us now prove the converse implication. Assume that $\mathfrak{s}$ is a übereven cluster, and let us denote by $\mathfrak{c}_1, \ldots, \mathfrak{c}_{N-1}$ its children (with $N \geq 3$); we remark that $d_+(\mathfrak{s}) = d_-(\mathfrak{c}_i)$ for all i. Letting $D = D_{\mathfrak{s},d_+(\mathfrak{s})}$, we have that the elements of $\mathcal{R} \cup \{\infty\}$ reduce in $(\mathcal{X}_D)_s$ to the

N distinct points, $\overline{\mathfrak{c}_1}, \ldots, \overline{\mathfrak{c}_{N-1}}, \infty = \overline{\mathcal{R} \smallsetminus \mathfrak{s}}$. Now, since all of the $\mathfrak{c}_i$'s are viable, we have $b_+(\mathfrak{s}) = d_+(\mathfrak{s})$ by Lemma 8.8; meanwhile, since $\mathfrak{s}$ is also assumed to be viable, we have $b_-(\mathfrak{s}) < b_+(\mathfrak{s})$. We conclude, in particular, that the disc D is valid (see Theorem 6.18).

As before, from the fact that $D = D_{\mathfrak{s},d_+(\mathfrak{s})}$ is a valid disc, we deduce that $b_+(\mathfrak{s}) = d_+(\mathfrak{s}) = b_-(\mathfrak{c}_i) = d_-(\mathfrak{c}_i)$. Since both $\mathfrak{s}$ and the $\mathfrak{c}_i$'s are viable, they have even cardinality and satisfy $b_-(\mathfrak{c}_i) < b_+(\mathfrak{c}_i)$ and $b_-(\mathfrak{s}) < b_+(\mathfrak{s})$, hence, by Proposition 6.17(a) we have that $\ell(\mathcal{X}_D, \overline{\mathfrak{c}_i}) = 0$ for all i, and $\ell(\mathcal{X}_D, \infty) = 0$. But this means that $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ is an étale double cover of the line, i.e. that $(\mathcal{Y}_D)_s$ is reducible, as we discussed in §4.5. $\square$

Proof of Theorem 8.1. The toric rank of a semistable k-curve is just the number of nodes (which we denote by N_{nodes}) minus the number of irreducible components (that we denote N_{irr}) plus 1 (see §2.1.1.7). Now, since we have $g(X) = 0$, the toric rank of $(\mathcal{X}^{(\mathrm{rst})})_s$ is 0 (as is the toric rank of the special fiber of any model of the line). Hence, the toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$ can be computed as

$$[N_{\mathrm{nodes}}((\mathcal{Y}^{\mathrm{rst}})_s) - N_{\mathrm{nodes}}((\mathcal{X}^{(\mathrm{rst})})_s)] - [N_{\mathrm{irr}}((\mathcal{Y}^{\mathrm{rst}})_s) - N_{\mathrm{irr}}((\mathcal{X}^{(\mathrm{rst})})_s)].$$

Now it follows from Proposition 8.4 that the first difference equals the number of viable clusters; meanwhile, Proposition 8.9 implies the second difference equals the number of übereven clusters. $\square$

2. The abelian rank

In this subsection we show how to compute the abelian rank of $(\mathcal{Y}_D)_s$ for any valid disc D, provided that, for each cluster $\mathfrak{s}$ linked to D, we are able to compute the invariants $b_\pm(\mathfrak{s})$ and $\lambda_\pm(\mathfrak{s})$ introduced in §6.2.

Proposition 8.10. Let $D := D_{\alpha,b}$ be a valid disc. In the $p \neq 2$ setting, then the genus of $\widetilde{(\mathcal{Y}_D)_s}$ equals $-1 + N_{\mathrm{odd}}/2$, where N_{odd} is the number of odd-cardinality clusters to which D is linked. In the $p = 2$ setting, instead we have the following.

(a) If D is linked to no cluster, then $g(\widetilde{(\mathcal{Y}_D)_s}) = -1 + (1 + \lambda_-(\varnothing, \alpha))/2$.
(b) If D is linked to a unique cluster $\mathfrak{s}$, then one of the three possibilities below holds:
  (i) $b = b_-(\mathfrak{s}) < b_+(\mathfrak{s})$, in which case $g(\widetilde{(\mathcal{Y}_D)_s}) = -1 + (1 + \lambda_-(\mathfrak{s}))/2$;
  (ii) $b = b_+(\mathfrak{s}) > b_-(\mathfrak{s})$, in which case $g(\widetilde{(\mathcal{Y}_D)_s}) = -1 + (1 + \lambda_+(\mathfrak{s}))/2$; or
  (iii) $b = b_+(\mathfrak{s}) = b_-(\mathfrak{s})$, in which case $g(\widetilde{(\mathcal{Y}_D)_s}) = -1 + (1 + \lambda_-(\mathfrak{s}))/2 + (1 + \lambda_+(\mathfrak{s}))/2$.
(c) If D is linked to $N \geq 3$ clusters, i.e. to a cluster $\mathfrak{s}_N$ and all of its children $\mathfrak{s}_1, \ldots, \mathfrak{s}_{N-1}$, then

$$g(\widetilde{(\mathcal{Y}_D)_s}) = -1 + \varepsilon_N(1 + \lambda_-(\mathfrak{s}_N))/2 + \sum_{i=1}^{N-1} \varepsilon_i(1 + \lambda_+(\mathfrak{s}_i))/2,$$

where $\varepsilon_i \in \{0, 1\}$, and ε_i is 1 (resp. 0) when $\mathfrak{s}_i$ is not viable (resp. viable), which in this setting is equivalent to the condition $b_-(\mathfrak{s}_i) = b_+(\mathfrak{s}_i)$ (resp. $b_-(\mathfrak{s}_i) < b_+(\mathfrak{s}_i)$).

Proof. The separable cover $\widetilde{(\mathcal{Y}_D)_s} \to (\mathcal{X}_D)_s$ is only ramified above the N points $P_1, \ldots, P_N = \infty$ of $(\mathcal{X}_D)_s$ to which the elements of $\mathcal{R} \cup \{\infty\}$ reduce, and we will apply the formula given in Proposition 4.28 to compute the genus of $\widetilde{(\mathcal{Y}_D)_s}$ based on the index $\ell(\mathcal{X}_D, P_i)$ defined there.

If $N = 1$ (which can happen only if $p = 2$), then D is linked to no cluster; Theorem 6.18 then ensures that $b = b_-(\varnothing, \alpha)$, and $\ell(\mathcal{X}_D, P_1) = \ell(\mathcal{X}_D, \infty) = \lambda_-(\varnothing, \alpha) + 1$ by Proposition 6.17(c).

If $N = 2$ (which can happen only if $p = 2$), then $\mathfrak{s} = D \cap \mathcal{R}$ is the unique cluster to which D is linked, and this implies that b is an internal point of $I(\mathfrak{s})$, i.e. $d_-(\mathfrak{s}) < b < d_+(\mathfrak{s})$; the 2 points of $(\mathcal{X}_D)_s$ to which the elements of $\mathcal{R} \cup \{\infty\}$ reduce are $P_1 = 0$ and $P_2 = \infty$. If we have $b = b_-(\mathfrak{s}) = b_+(\mathfrak{s})$, then parts (c) and (d) of Proposition 6.17 give $\ell(\mathcal{X}_D, P_1) = 1 + \lambda_+(\mathfrak{s})$ and $\ell(\mathcal{X}_D, P_2) = 1 + \lambda_-(\mathfrak{s})$ respectively. We now assume that $b_-(\mathfrak{s}) < b_+(\mathfrak{s})$. If we have $b = b_-(\mathfrak{s}) < b_+(\mathfrak{s})$, then parts (a) and (c) of Proposition 6.17 give $\ell(\mathcal{X}, P_1) = 0$ and $\ell(\mathcal{X}, P_2) = 1 + \lambda_-(\mathfrak{s})$ respectively, while, if instead we have $b = b_+(\mathfrak{s}) > b_-(\mathfrak{s})$, points (a) and (d) of Proposition 6.17 give $\ell(\mathcal{X}, P_1) = 1 + \lambda_+(\mathfrak{s})$ and $\ell(\mathcal{X}, P_2) = 0$ respectively.

Now suppose that we have $N \geq 3$; this means that D is linked to a cluster $\mathfrak{s}_N$ and to all of its children $\mathfrak{s}_1, \ldots, \mathfrak{s}_{N-1}$, which must all have even cardinality if $p = 2$. We have $b = d_+(\mathfrak{s}_N) = d_-(\mathfrak{s}_i)$ for all $i = 1, \ldots, N-1$; by Theorem 6.18(a), we moreover know that, since D is a valid disc, we must have $b = b_+(\mathfrak{s}_N) = b_-(\mathfrak{s}_i)$; this also forces $b_-(\mathfrak{s}_i) \leq b_+(\mathfrak{s}_i)$ for $1 \leq i \leq N$. Now choose any $i \in \{1, \ldots, N-1\}$. If we have $b_+(\mathfrak{s}_i) = b_-(\mathfrak{s}_i)$ (which can only occur if $p = 2$), then we obtain from Proposition 6.17(d) that $\ell(\mathcal{X}_D, P_i) = 1 + \lambda_+(\mathfrak{s}_i)$, where P_i is the point to which the roots of $\mathfrak{s}_i$ reduce. If on the other hand we have $b_+(\mathfrak{s}_i) > b_-(\mathfrak{s}_i)$, then we obtain from Proposition 6.17(a),(b) that $\ell(\mathcal{X}_D, P_i) = 0$ (resp. $\ell(\mathcal{X}_D, P_i) = 1$) if the cardinality $|\mathfrak{s}_i|$ is even (resp. odd) (it is always even if $p = 2$). Similarly, using Proposition 6.17(a),(b),(d), we obtain that $\ell(\mathcal{X}_D, P_N) = 1 + \lambda_-(\mathfrak{s})$ when $b_+(\mathfrak{s}_N) = b_-(\mathfrak{s}_N)$, while, if $b_+(\mathfrak{s}_N) > b_-(\mathfrak{s}_N)$, we have $\ell(\mathcal{X}_D, P_N) = 0$ (resp. $\ell(\mathcal{X}_D, P_N) = 1$) if the cardinality $|\mathfrak{s}_N|$ is even (resp. odd).

Now the claimed formulas for $g(\widetilde{(\mathcal{Y}_D)_s})$ follows directly from applying Proposition 4.28. ☐

Remark 8.11. The genus $g(\widetilde{(\mathcal{Y}_D)_s})$ coincides with the abelian rank of $(\mathcal{Y}_D)_s$, unless $(\mathcal{Y}_D)_s$ consists of 2 components, in which case we have $g(\widetilde{(\mathcal{Y}_D)_s}) = -1$ while the abelian rank of $(\mathcal{Y}_D)_s$ is 0.

Corollary 8.12. Suppose that D is a valid disc that is linked to no cluster, or that it is linked to only one cluster $\mathfrak{s}$ and is the unique valid disc linked to $\mathfrak{s}$. Then the k-curve $(\mathcal{Y}_D)_s$ is irreducible and has abelian rank ≥ 1.

Proof. This follows immediately from Proposition 8.10(a) and Proposition 8.10(b)(iii), taking into account that, when D is linked to no cluster, we have $\lambda_-(\varnothing, \alpha) \geq 3$ by Theorem 6.18. ☐

3. Partitioning the components of the special fiber

Suppose that in the $p = 2$ setting we are given a disc $D := D_{\alpha,b}$ with $\alpha \in \bar{K}$ and $b \in \mathbb{Q}$ and that the cover $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable, i.e. D satisfies $\mathfrak{t}^{\mathcal{R}}(D) < 2v(2)$, and let $\beta \in \bar{K}^\times$ be such that $v(\beta) = b$. We have seen in §4.6 that, in this case, given a part-square decomposition $f = q^2 + \rho$ that is good at the disc D, the special fiber $(\mathcal{Y}_D)_s$ is described by an equation of the form $y^2 = \overline{\rho_0}(x_{\alpha,\beta})$, where $\overline{\rho_0}$ is a normalized reduction of $\rho_{\alpha,\beta}$. Let us also recall Definition 4.33 which says that, given $P \in (\mathcal{X}_D)_s$, we denote by $\mu(\mathcal{X}_D, P)$ the order of vanishing of $\overline{\rho_0}'$ at P, which is an even integer.

Letting $\mathfrak{D}^{(\mathrm{rst})}$ denote the collection of discs corresponding to $\mathcal{X}^{(\mathrm{rst})}$ (in the sense of §4.2), for any $P \in (\mathcal{X}_D)_s$ we write $\mathfrak{D}_P \subseteq \mathfrak{D}^{(\mathrm{rst})}$ for the non-empty subset consisting of those $D' \in \mathfrak{D}^{(\mathrm{rst})}$ such that $\mathrm{Ctr}(\mathcal{X}_D, \mathcal{X}_{D'}) = \{P\}$.

Proposition 8.13. In the setting above, the following are equivalent:

(a) $\mu(\mathcal{X}_D, P) > 0$;
(b) $(\mathcal{Y}_D)_s$ is singular above P;
(c) $\mathfrak{D}_P \neq \varnothing$;
(d) $\mathfrak{D}_P$ contains a valid disc.

Proof. The equivalence between (a) and (b) was already discussed in §4.6, while the equivalence between (a)/(b) and (c) is an immediate consequence of Corollary 3.20. Finally, in light of Proposition 4.37, it is easy to see that $\mathcal{D}_P$ contains a valid disc whenever it is non-empty. $\square$

As in §4.6, we let R_{sing} denote the set of points of $(\mathcal{X}_D)_s$ satisfying the equivalent conditions above. Given any $P \in R_{\mathrm{sing}}$, we write $\mathcal{X}_P^{(\mathrm{rst})}$ for the model of the line X corresponding to $\mathfrak{D}_P$ and let $\mathcal{Y}_P^{\mathrm{rst}}$ be the corresponding model of Y. We now present the main result of this subsection.

Proposition 8.14. In the setting above, the model $\mathcal{Y}_P^{\mathrm{rst}}$ satisfies the following properties:

(a) the strict transform C_P of $(\mathcal{Y}_P^{\mathrm{rst}})_s$ in $(\mathcal{Y}^{\mathrm{rst}})_s$ intersects the rest of $(\mathcal{Y}^{\mathrm{rst}})_s$ at a single node (when it does not coincide with the whole special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$); and
(b) the arithmetic genus of the k-curve C_P (which is to say, the sum of the abelian and toric rank of the special fiber of $\mathcal{Y}_P^{\mathrm{rst}}$) is equal to $\frac{1}{2}\mu(\mathcal{X}_D, P)$.

Remark 8.15. An analogous result holds in the $p \neq 2$ setting if D is taken to be any disc and $P \in (\mathcal{X}_D)_s(k)$ is a point over which $(\mathcal{Y}_D)_s$ exhibits a unibranch singularity (i.e. $P \in R_1$ in the language of §4.5): in this situation the invariant $\mu(\mathcal{X}_D, P)$ is given by $N_P - 1$, where N_P is the (necessarily odd) number of roots of $\mathcal{R} \cup \{\infty\}$ reducing to P. The proof is analogous to (and in some aspects simpler than) that of Proposition 8.14.

Before presenting the proof, let us introduce the following two lemmas.

Lemma 8.16. Choose $\alpha \in \bar{K}$, and suppose that, for some rational number $b \in \mathbb{Q}$, we have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha,b}) < 2v(2)$. Let us fix a part-square decomposition $f = q^2 + \rho$ that is totally odd with respect to a center $\alpha \in \bar{K}$. Then the function $c \mapsto \underline{v}_\rho(D_{\alpha,c})$ is not differentiable at the input $c = b$ if and only if there is a valid disc $D_{\alpha',b'}$ such that $v(\alpha' - \alpha) = b$ and $b' > b$.

Proof. By Lemma 6.2(b) the function $c \mapsto \underline{v}_\rho(D_{\alpha,c})$ is differentiable at $c = b$ if and only if some (any) normalized reduction of $(\rho_{\alpha,\beta})'$, for β such that $v(\beta) = b$, has a root $P \in (\mathcal{X}_{D_{\alpha,\beta}})_s$ which is neither $\overline{x_{\alpha,b}} = 0$ nor $\overline{x_{\alpha,\beta}} = \infty$. By applying Proposition 8.13 to the disc $D_{\alpha,b}$ (taking into account Proposition 4.6), this is equivalent to saying that there exists a valid disc $D_{\alpha',b'}$ such that $v(\alpha' - \alpha) = b$ and $b' > b$. $\square$

Remark 8.17. It is clear from Lemma 6.2 and definitions of the functions involved that the non-differentiability condition in the statement of the above lemma is satisfied whenever there is an input $c = b$ which is not equal to the depth of any cluster containing α and at which the function $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha,c})$ is not differentiable.

Lemma 8.18. Let $D := D_{\alpha,b}$ be a disc such that $\mathfrak{t}^{\mathcal{R}}(D) < 2v(2)$, and let $D^{\varepsilon} = D_{\alpha,b-\varepsilon}$ for some $\varepsilon > 0$. Let $r : (\mathcal{X}_D)_s(k) \to (\mathcal{X}_{D^{\varepsilon}})_s(k)$ be the map taking all points $P \in (\mathcal{X}_D)_s(k) \smallsetminus \{\infty\}$ to 0 and taking $\infty \in (\mathcal{X}_D)_s(k)$ to $\infty \in (\mathcal{X}_{D^{\varepsilon}})_s(k)$. Then if ε is sufficiently small, we have $\mu(\mathcal{X}_{D^{\varepsilon}}, Q) = \sum_{P \in r^{-1}(Q)} \mu(\mathcal{X}_D, P)$ for all $Q \in (\mathcal{X}_{D^{\varepsilon}})_s(k)$, and moreover, we have the inclusions

$$\bigsqcup_{P \in r^{-1}(Q)} \mathfrak{D}_P \subseteq \mathfrak{D}_Q \subseteq \Big(\bigsqcup_{P \in r^{-1}(Q)} \mathfrak{D}_P \Big) \cup \{D\}. \tag{45}$$

Proof. It is easy to see that both claims of the lemma hold when ε is chosen such that there is no disc $D_{\alpha',b'} \in \mathfrak{D}^{(\mathrm{rst})}$ with $b - \varepsilon < v(\alpha' - \alpha) < b$; the second (resp. first) inequality in (45) is an equality if D is (resp. is not) a disc in $\mathfrak{D}^{(\mathrm{rst})}$. $\square$

Proof (of Proposition 8.14). Let $D, P, q, \rho, \mathfrak{D}^{(\mathrm{rst})}$ and $\mathfrak{D}_P$ be as specified at the beginning of this subsection; in particular, since $\mathfrak{t}^{\mathcal{R}}(D) < 2v(2)$, we are in the $p = 2$ setting. Let us write $D = D_{\alpha,b}$ for some $\alpha \in \bar{K}$ and $b \in \mathbb{Q}$, and let $D_{\varepsilon} = D_{\alpha,b+\varepsilon}$ for ε arbitrarily small. In this proof, we will always assume for simplicity that $\overline{x_{\alpha,\beta}}(P) = 0$ (which is certainly the case for an appropriate choice of α); according to Proposition 4.6 this means that, given a disc D', we have $D' \in \mathfrak{D}_P$ if and only if we have $D' \in \mathfrak{D}^{(\mathrm{rst})}$ and $D' \subseteq D_{\varepsilon}$.

Let us first address part (a). Let P' denote the node of $(\mathcal{X}^{(\mathrm{rst})})_s$ at which the strict transform of $(\mathcal{X}_P^{(\mathrm{rst})})_s$ intersects the rest of $(\mathcal{X}^{(\mathrm{rst})})_s$, and let $D_1 \in \mathfrak{D}_P$ and $D_2 \in \mathfrak{D}^{(\mathrm{rst})} \smallsetminus \mathfrak{D}_P$ be the discs corresponding to the two lines of $(\mathcal{X}^{(\mathrm{rst})})_s$ meeting at P'. Suppose by way of contradiction that P' has two distinct inverse images in $(\mathcal{Y}^{\mathrm{rst}})_s$; then Proposition 8.4 ensures that D_1 and D_2 are the 2 valid discs linked to some viable cluster $\mathfrak{s}$ and so in particular are not disjoint. From this, since $D_1 \subseteq D_{\varepsilon}$ but $D_2 \not\subseteq D_{\varepsilon}$, it follows that we have $D_1 \subseteq D_{\varepsilon} \subseteq D_2$ for all small enough ε. This means that D itself is linked to $\mathfrak{s}$ and its depth b lies in the interval $J(\mathfrak{s})$, and hence we have $\mathfrak{t}^{\mathcal{R}}(D) = 2v(2)$, which contradicts our assumption that $(\mathcal{Y}_D)_s \to (\mathcal{X}_D)_s$ is inseparable. We conclude that P' must have a single inverse image in $(\mathcal{Y}^{\mathrm{rst}})_s$, and (a) is proved.

Let us now address part (b). We write $\mu(\mathcal{X}_D, P) = 2\nu$ where ν is a positive integer, and we let $D'_1, \ldots, D'_h$ be the maximal valid discs that are contained in D_{ε}; for each i, let us moreover choose $\alpha'_i \in \bar{K}$ such that $D'_i = D_{\alpha'_i,b'_i}$ with $b_i \in \mathbb{Q}$. Since the discs D'_i are valid, we have $\mathfrak{t}^{\mathcal{R}}(D'_i) = 2v(2)$; thanks to the maximality assumption, we can be sure, in light of Theorem 6.18, that we have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha'_i,c}) < 2v(2)$ for all $c \in [b, b'_i)$.

We proceed by induction on $n := \max_{1 \leq i \leq h}(|D'_i \cap \mathcal{R}|)$, beginning by proving the result for $n = 0$. We preliminarily observe that, for a fixed $n \geq 0$, it is enough to address the case where $h = 1$. Indeed, one can verify, by repeatedly applying Lemma 8.18, that the result of (b) is true for the disc D and the point $P \in (\mathcal{X}_D)_s$ if it is true for the discs $D_{\alpha'_i,b'_i-\varepsilon}$ at the points $\overline{\alpha'_i} \in (\mathcal{X}_{D_{\alpha'_i,b'_i-\varepsilon}})_s(k)$.

Assume that $h = 1$; for simplicity of notation we write $D' := D_{\alpha',b'}$ for D'_1. We moreover write $\mathfrak{s}' = D' \cap \mathcal{R}$ and choose a part-square decomposition $f = \tilde{q}^2 + \tilde{\rho}$ that is totally odd with respect to the center α'. Since we have $\mathfrak{t}^{\mathcal{R}}(D_{\alpha',c}) < 2v(2)$ for all $c \in [b, b')$ and D' is a valid disc, we have $b' = b_-(\mathfrak{s}', \alpha')$. Let us moreover remark that, by Lemma 6.2, the right derivative of $c \mapsto \underline{v}_{\tilde{\rho}}(D_{\alpha',c})$ at $c = b$ is equal to $2\nu + 1$ and, since D' is the maximum among the valid discs contained in D_ε, the slope of the function $c \mapsto \underline{v}_{\tilde{\rho}}(D_{\alpha',c})$ actually remains equal to $2\nu + 1$ for all $c \in [b, b']$ by Lemma 8.16. In particular, by applying Lemma 6.2 to $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha',c}) = \underline{v}_{\tilde{\rho}}(c) - \underline{v}_{f_{\alpha',1}}(c)$, we can be sure that $\lambda_-(\mathfrak{s}', \alpha') = 2\nu + 1 - n$, recalling that $n = |\mathfrak{s}'|$ in this situation.

We now set out to prove the case $n = 0$ and $h = 1$. In this case, the only valid disc contained in D_ε is D', and we have $\mathfrak{D}_P = \{D'\}$ and that $(\mathcal{Y}_P^{\mathrm{rst}})_s$ has toric rank 0 and abelian rank $-1 + (\lambda_-(\varnothing, \alpha') + 1)/2 = \nu$ by Proposition 8.10(a), as we wanted.

Now assume that $n \geq 2$ and $h = 1$, and assume inductively that the conclusion of part (b) holds for all lesser values of n for all h. We can clearly write $\mathfrak{s}' = \mathfrak{s}_1 \sqcup \ldots \sqcup \mathfrak{s}_r$, where the $\mathfrak{s}_i$'s are the (even-cardinality) clusters contained in $\mathfrak{s}'$ such that $b_-(\mathfrak{s}_i) \leq b_+(\mathfrak{s}_i) < d_+(\mathfrak{s}_i)$ and such that we have $\mathfrak{t}^{\mathcal{R}}(D_{\mathfrak{s}_i,c}) = 2v(2)$ for $c \in [b, b_+(\mathfrak{s}_i)]$ and $\mathfrak{t}^{\mathcal{R}}(D_{\mathfrak{s}_i,c}) < 2v(2)$ for $c \in (b_+(\mathfrak{s}_i), d_+(\mathfrak{s}_i)]$. For each i, we write the following: let $f^2 = q_i^2 + \rho_i$ be a part-square decomposition that is totally odd with respect to α_i; let ν_i be the integer such that the right derivative of $c \mapsto \underline{v}_{\rho_i}(D_{\mathfrak{s}_i,c})$ at $c = b_+(\mathfrak{s}_i)$ equals $2\nu_i + 1$; and denote the disc $D_{\mathfrak{s}_i, b_+(\mathfrak{s}_i)+\varepsilon}$ by D_i. Note that we have the formula $\lambda_+(\mathfrak{s}_i) = |\mathfrak{s}_i| - (2\nu_i + 1)$.

Let $\mathcal{J} \subseteq \mathcal{I} := \{1, \ldots, r\}$ be the subset of indices such that $d_-(\mathfrak{s}_i) = b_-(\mathfrak{s}_i) = b_+(\mathfrak{s}_i)$, i.e. the indices for which $\mathfrak{s}_i$ is not viable; we define the partition $\mathcal{J} = \mathcal{J}_1 \sqcup \ldots \sqcup \mathcal{J}_s$ so that each $\mathcal{J}_j \subseteq \mathcal{J}$ is a maximal subset of indices corresponding to sibling clusters. We classify the valid discs contained in D as follows.

(I) For each $i \in \mathcal{I}$, we have the valid discs contained in D_i; by the inductive hypothesis, these give a total contribution of $\sum_{1 \leq i \leq r} \nu_i$ to the sum of the abelian and toric ranks of $(\mathcal{Y}_P^{\mathrm{rst}})_s$.

(II) We have the discs $D_{\mathfrak{s}_i, b_+(\mathfrak{s}_i)}$ for all $i \in \mathcal{I} \smallsetminus \mathcal{J}$, along with the disc D' when $\mathfrak{s}'$ is viable. Now we observe that
  (a) each of the discs $D_{\mathfrak{s}_i, b_+(\mathfrak{s}_i)}$ contributes a component of $(\mathcal{Y}_P^{\mathrm{rst}})_s$ of abelian rank equal to $-1 + (1 + \lambda_+(\mathfrak{s}_i))/2 = -1 + \frac{1}{2}|\mathfrak{s}_i| - \nu_i$ by Proposition 8.10(b)(ii); and
  (b) when $\mathfrak{s}'$ is viable, the disc D' contributes a component of $(\mathcal{Y}_P^{\mathrm{rst}})_s$ of abelian rank equal to $-1 + (1 + \lambda_-(\mathfrak{s}'))/2 = \nu - \frac{1}{2}|\mathfrak{s}'|$ by Proposition 8.10(b)(i).

(III) We have the valid discs $D_{\mathfrak{s}_i, d_-(\mathfrak{s}_i)}$ for $i \in \mathcal{J}$, which are precisely the (distinct) discs $D_{\mathfrak{r}_j, d_+(\mathfrak{r}_j)}$ for $1 \leq j \leq s$ where each $\mathfrak{r}_j$ is the common parent of the clusters $\mathfrak{s}_i$ with $i \in \mathcal{J}_j$ and $\tilde{\alpha}_j = \alpha_i$ for a choice of $i \in \mathcal{J}_j$. The abelian rank of $(\mathcal{Y}_{D_{\mathfrak{r}_j, d_+(\mathfrak{r}_j)}})_s$ can be computed using Proposition 8.10(c) as follows:
  (a) if $\mathfrak{r}_j \subsetneq \mathfrak{s}'$ or if $\mathfrak{r}_j = \mathfrak{s}'$ and $\mathfrak{s}'$ is viable, it is equal to
$$-1 + \sum_{i \in \mathcal{J}_j} \frac{1}{2}(1 + \lambda_+(\mathfrak{s}_i)) = -1 + \sum_{i \in \mathcal{J}_j} (\frac{1}{2}|\mathfrak{s}_i| - \nu_i); \text{ and}$$
  (b) if $\mathfrak{r}_j = \mathfrak{s}'$ and $\mathfrak{s}'$ is not viable, then it is equal to
$$-1 + \frac{1}{2}(1 + \lambda_-(\mathfrak{s}')) + \sum_{i \in \mathcal{J}_j} \frac{1}{2}(1 + \lambda_+(\mathfrak{s}_i))) = \sum_{i \in \mathcal{J}_j} (\frac{1}{2}|\mathfrak{s}_i| - \nu_i) + (\nu - \frac{1}{2}|\mathfrak{s}'|).$$

Moreover, if $\mathfrak{s}'$ is not viable but there is no index $j \in \{1, \ldots, s\}$ such that $\mathfrak{r}_j = \mathfrak{s}'$, we include D' in the subset of discs of type III(b); the disc D' contributes $\nu - \frac{1}{2}|\mathfrak{s}'|$ to the abelian rank by Proposition 8.10(c).

(IV) All other valid discs contained in D are of the form $D_{\mathfrak{r},d_+(\mathfrak{r})}$ for $\mathfrak{r} \subsetneq \mathfrak{s}$ a übereven cluster containing one of the clusters $\mathfrak{s}_i$; these each contribute 2 lines to $(\mathcal{Y}_P^{\mathrm{rst}})_s$ and thus do not increase the abelian rank.

The discs of type II, III and IV give a contribution to the abelian rank of C_P that adds up to $\nu - \sum_{1 \leq i \leq r} \nu_i - t$, where t is the number of valid discs of type II(a) and III(a). Their contribution to the toric rank of C_P equals t by Theorem 8.1, taking into account that t equals the number of viable non-übereven clusters $\mathfrak{r} \subseteq \mathfrak{s}'$ such that $\mathfrak{s}_i \subsetneq \mathfrak{r}$ for some $i \in \mathcal{I}$. Thus, the valid discs contained in D give a total contribution to the abelian and toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$ equal to ν, which is what we wanted. □

Corollary 8.19. Suppose that $\mathfrak{s}$ is a cluster of odd cardinality $2\nu+1$ with $1 \leq \nu \leq g-1$. Then the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of two G-invariant k-curves C_0 and C_∞ meeting as a single node $Q_{\mathfrak{s}} \in (\mathcal{Y}^{\mathrm{rst}})_s$; their arithmetic genera are ν and $g - \nu$ respectively.

Proof. Choose D to be any disc of the form $D_{\mathfrak{s},b}$ for some $b \in (d_-(\mathfrak{s}), d_+(\mathfrak{s}))$, and let $\alpha \in D_{\mathfrak{s},b}$ be a center and $\beta \in \bar{K}^\times$ be an element of valuation b. We have that ν roots of $\mathcal{R} \cup \{\infty\}$ reduce to $\overline{x_{\alpha,\beta}} = 0$, while the remaining $2g - \nu$ roots reduce to $\overline{x_{\alpha,\beta}} = \infty$ in $(\mathcal{X}_D)_s$.

Assume that we are in the $p = 2$ setting. The normalized reduction of $\rho_{\alpha,\beta}$ has the form $x_{\alpha,\beta}^{2m+1}$; meanwhile, the part-square decomposition $f = 0^2 + f$ is good at D by Proposition 4.18. We deduce in particular that $\mathfrak{t}^{\mathcal{R}}(D) = 0$, and that we have that $\mu(\mathcal{X}_D, 0) = 2\nu$ and $\mu(\mathcal{X}_D, \infty) = 2g - 2\nu$ and that $\mu(\mathcal{X}_D, P) = 0$ at all other points $P \in (\mathcal{X}_D)_s$ (see §4.6). As a consequence of Proposition 4.35, we have $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$. Now the corollary follows as an immediate application of Proposition 8.14. In the $p \neq 2$ setting, we also have $\mathcal{X}_D \not\leq \mathcal{X}^{(\mathrm{rst})}$ by Theorem 4.32, and the corollary follows from Remark 8.15. □

CHAPTER 9

Computations for hyperelliptic curves of low genus

In this section we apply our results from §6,7,8 to determine the possible structures of special fibers of relatively stable models $\mathcal{Y}^{\mathrm{rst}}$ of hyperelliptic curves Y defined by equations of the form $y^2 = f(x)$ (with $\deg(f) = 2g + 1$) over residue characteristic $p = 2$, given the cluster data associated to f along with (when the genus $g \geq 2$) the valuations of elements of K coming from formulas involving the coefficients of certain factors of f. For genera $g = 1, 2$, we shall classify hyperelliptic curves over K into several cases that depend on the aforementioned data and show how to compute each component of $(\mathcal{Y}^{\mathrm{rst}})_s$ along with its toric rank on such a case-by-case basis.

In order to simplify notation in the formulas and conditions appearing in our statements below, the hypotheses of each of the results of this section we include simplifying assumptions that

(1) f is monic;
(2) the depth of the full set of roots $\mathcal{R}$ is 0; and
(3) one of the roots of f (namely one which is contained in a particular even-cardinality cluster $\mathfrak{s}$ we are working with) is 0.

Assumption (1) holds after appropriately scaling the y-coordinate (by a scalar which lies in at most a quadratic extension of K). Assumptions (2) and (3) hold after making simple changes of coordinates of the defining equation of the hyperelliptic curve which translate and scale the roots of f; this is done by translating and scaling the x-coordinate and again appropriately scaling the y-coordinate.

Remark 9.1. Assume that f is monic and one of its roots is 0, i.e. it satisfies (1) and (3). Then, condition (2) just means that f has integral coefficients, and at least one of its non-leading coefficients is a unit; equivalently, f has integral roots, and one of its roots is a unit.

The following lemma will help us below to characterize those valid discs which are not linked to any cluster.

Lemma 9.2. Suppose f is monic; let $\alpha \notin \mathcal{R}$, and let $I(\varnothing, \alpha) = [d_-(\varnothing, \alpha), +\infty)$ be the interval defined in Definition 6.9. Then we have $d_-(\varnothing, \alpha) = \max_{a \in \mathcal{R}} v(a - \alpha)$, and, for all discs $D := D_{\alpha,b}$ with $b \in I(\varnothing, \alpha)$, we have

$$\underline{v}_f(D) = \sum_{a \in \mathcal{R}} v(a - \alpha). \tag{46}$$

Proof. This is an immediate computation. $\square$

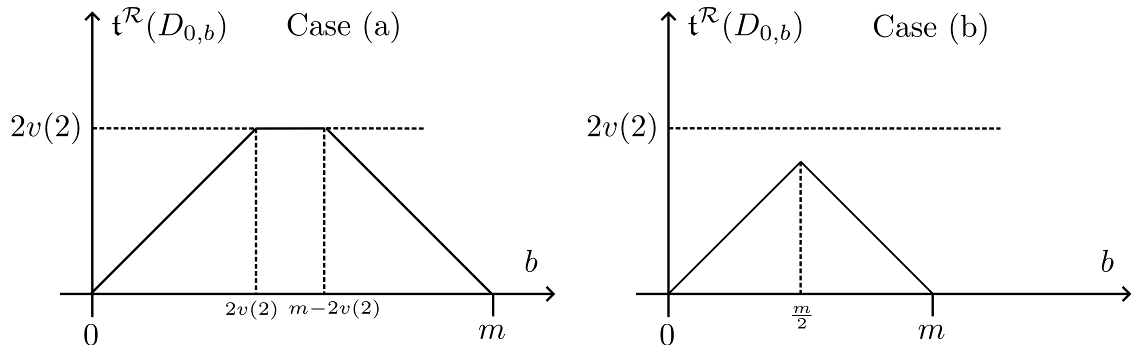

FIGURE 1. The shape of the function $I(\mathfrak{s}) \to [0, 2v(2)]$, $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ in cases (a) and (b) of Theorem 9.3 provided that $m > 0$.

1. The $g = 1$ case (elliptic curves)

We begin our search for concrete results for hyperelliptic curves by considering the simplest situation: the case that $g = 1$ so that Y is an elliptic curve. Suppose that $Y : y^2 = f(x)$ is an elliptic curve over K, i.e. we have $g = 1$ and $\deg(f) = 3$. We note that this case is treated (in a more concrete and elementary fashion) as the main topic of J. Yelton's paper [17]. We label the three roots of f as $a_1 := 0, a_2, a_3 \in \bar{K}$. Apart from the full set of roots, clearly the only non-singleton cluster we may have is a cluster of cardinality 2 which we assume coincides with $\{a_1 = 0, a_2\}$. The J. Yelton's previous results [17, Theorems 1 and 4] may be rephrased using the terminology of this work and adapted to our particular desired semistable model $\mathcal{Y}^{\mathrm{rst}}$ as follows.

THEOREM 9.3. With the above set-up, let $m = v(a_2)$ (so that $m = 0$ if and only if there is no cardinality-2 cluster and otherwise m is the depth of the cardinality-2 cluster $\mathfrak{s} = \{0, a_2\}$).

(a) Suppose that $m > 4v(2)$. Then there are exactly 2 valid discs $D_+ := D_{0,m-2v(2)}$ and $D_- := D_{0,2v(2)}$, both linked to the cluster $\mathfrak{s}$. Thus, the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of 2 components each of abelian rank 0 which intersect at 2 points.

(b) Suppose that $m \leq 4v(2)$. Then there is exactly 1 valid disc D_{α_1,b_1}, where α_1 satisfies $v(\alpha_1) = v(\alpha_1 - a_2) = \frac{1}{2}m$ and $v(\alpha_1 - a_3) = 0$, while $b_1 = \frac{1}{3}(m + 2v(2))$; moreover, the center $\alpha_1 \in \bar{K}$ can be taken to be a root of the polynomial
$$F(T) := P_1^2(T) - 4P_2(T)P_0(T) \in K[T],$$
where each $P_i(T) \in K[T]$ is the z^i-coefficient of $f(z + T) \in K[T][z]$. The corresponding model $\mathcal{Y}_D$ of Y has smooth special fiber; thus, in this case, the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$ coincides with $\mathcal{Y}_D$ and Y attains good reduction.

Remark 9.4. The cases (a) and (b) of Theorem 9.3 (when $m > 0$) correspond to the possible shapes of the function $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ as b ranges in $I(\mathfrak{s}) = [d_-(\mathfrak{s}), d_+(\mathfrak{s})] = [0, m]$, which are described in Figure 1.

Remark 9.5. The case that treated in [17] is when the elliptic curve is a member of the Legendre family, i.e. of the form $E_\lambda : y^2 = f(x) := x(x-1)(x-\lambda)$ for some $\lambda \in K \smallsetminus \{0,1\}$ with $m = v(\lambda)$ (in other words, we set $a_2 = \lambda$ and $a_3 = 1$, which can always done after appropriate translation and scaling). In this situation, we make the following observations.

(a) The above theorem directly implies that the elliptic curve E_λ has potentially good reduction if and only if $m \leq 4v(2)$. This could alternately be deduced as a consequence of the following facts. It is well known (see for instance [14, §IV.1.2] or [16, Proposition VII.5.5]) that any elliptic curve over a complete discrete valuation field has good (resp. multiplicative) reduction over some finite extension of that field if and only if the valuation of its j-invariant is nonnegative (resp. negative). The formula for the j-invariant of the Legendre curve E_λ is given as in [16, Proposition III.1.7] by

$$j(E_\lambda) = 2^8 \frac{(\lambda^2 - \lambda + 1)^3}{\lambda^2(\lambda - 1)^2}. \tag{47}$$

It is easily computed from this formula that we have $v(j(E)) = 8v(2) - 2v(\lambda)$, and our claim about potentially good reduction if and only if $m \leq 4v(2)$ follows.

(b) We compute the formulas $P_2(T) = 3T - (\lambda+1)$, $P_1(T) = 3T^2 - 2(\lambda+1)T + \lambda$, and $P_0(T) = T^3 - (\lambda+1)T^2 + \lambda T$. Then the polynomial F given in Theorem 9.3(b) can be written in the simpler form

$$F(T) = -3T^4 + 4(1+\lambda)T^3 - 6\lambda T^2 + \lambda^2. \tag{48}$$

(c) For the polynomial defining E_λ, we have $f_+^{\mathfrak{s}}(z) = f_-^{\mathfrak{s}}(z) = 1 - z$, and the obvious totally odd part-square decompositions for both of them induce (as in §6.3.3.2) the decomposition $f(x) = [\sqrt{-1}x]^2 + [x^3 - \lambda x^2 + \lambda x]$ (for some choice of square root of -1), which according to Proposition 6.42 is good at the discs $D_{0,m-2v(2)}$ and $D_{0,2v(2)}$. This is helpful for explicitly constructing the components of $(E_\lambda^{\mathrm{rst}})_s$ in the case that $m \geq 4v(2)$.

Examples of computations which yield the desired model $\mathcal{Y}^{\mathrm{rst}}$ in the case that $m \leq 4v(2)$ are given as [17, Examples 2 and 3]. Below is an example for the $m > 4v(2)$ case, which is treated in [17, Example 9] except that there a semistable model whose special fiber has a single (nodal) component, rather than the relatively stable model $\mathcal{Y}^{\mathrm{rst}}$, is found.

Example 9.6. Let Y be the elliptic curve over $\mathbb{Z}_2^{\mathrm{unr}}$ given by

$$y^2 = x(x-64)(x-1),$$

so that we have a unique even-cardinality cluster $\mathfrak{s} = \{0, 64\}$ of relative (and absolute) depth $m = 6v(2)$. We are therefore in the situation of Theorem 9.3(a), and the valid discs can be taken to be $D_1 := D_- = D_{0,2v(2)}$ and $D_2 := D_+ = D_{0,4v(2)}$. Using the sufficiently odd part-square decomposition given by Remark 9.5(c), we obtain (see §4.4) that the changes in coordinates corresponding to each of these discs may be written as

$$x = 4x_1 = 16x_2, \quad y = 8y_1 + 4\sqrt{-1}x_1 = 32y_2 + 16\sqrt{-1}x_2.$$

We now get equations for the models $\mathcal{Y}_1$ and $\mathcal{Y}_2$ corresponding to D_1 and D_2 respectively as

$$\mathcal{Y}_1 : y_1^2 + \sqrt{-1}x_1y_1 = x_1^3 - 2^4x_1^2 + 2^2x_1, \qquad \mathcal{Y}_2 : y_2^2 + \sqrt{-1}x_2y_2 = 2^2x_2^3 - 2^4x_2^2 + x_2. \tag{49}$$

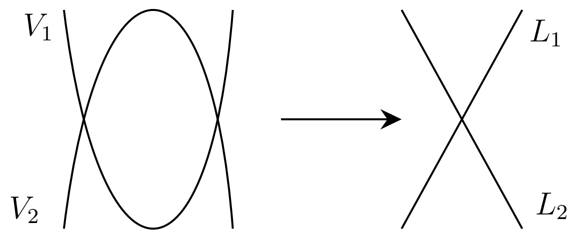

FIGURE 2. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$, shown above, mapping to $(\mathcal{X}^{(\mathrm{rst})})_s$; each component V_i of $(\mathcal{Y}^{\mathrm{rst}})_s$ maps to each component $L_i := (\mathcal{X}_{D_i})_s$ of $(\mathcal{X}^{(\mathrm{rst})})_s$.

whose special fibers are the $\bar{\mathbb{F}}_2$-curves described by the equations

$$(\mathcal{Y}_1)_s : y_1^2 + \sqrt{-1}x_1y_1 = x_1^3, \qquad (\mathcal{Y}_2)_s : y_2^2 + \sqrt{-1}x_2y_2 = x_2. \tag{50}$$

Note that $\mathcal{Y}_1$ is already a semistable model of Y, its special fiber being a curve with a node, and is one that could be obtained from [17, Theorem 4 and Remark 5], but it is not the relatively stable model as the node is a vanishing node (see Definition 3.8). The desingularizations of $(\mathcal{Y}_1)_s$ and $(\mathcal{Y}_2)_s$ are each smooth curves of genus 0 and give rise to the components of $(\mathcal{Y}^{\mathrm{rst}})_s$, the configuration of which is shown in Figure 2.

Throughout the rest of this subsection we prove Theorem 9.3. By Theorem 5.13(a), we know that for any valid disc D, we have either $D \cap \mathcal{R} = \varnothing$ or $D \cap \mathcal{R} = \mathfrak{s}$. Let us first treat the situation where $\mathfrak{s} := \{0, a_2\}$ is a cluster and search for all valid discs (if any) which contain it.

1.1. Finding valid discs containing a cardinality-2 cluster. For the moment, let us assume that $\mathfrak{s} := \{0, a_2\}$ is a cluster; we fix 0 as a center for any disc containing $\mathfrak{s}$. Then Proposition 6.28(a),(c) (along with Remark 6.29(c)) directly implies that we have $\mathfrak{t}_{\pm}^{\mathfrak{s}}(b) = \min\{b, 2v(2)\}$, so that $b_0(\mathfrak{t}_{\pm}^{\mathfrak{s}}) = 2v(2)$ and hence $B_{f,\mathfrak{s}} = 4v(2)$ (see also Corollary 1.6). From this, we deduce immediately that $(\mathcal{Y}^{\mathrm{rst}})_s$ has toric rank 1 and hence abelian rank 0 (resp. toric rank 0 and hence abelian rank 1) if and only if $m > 4v(2)$ (resp. $m \leq 4v(2)$): this is a consequence of Theorem 8.1. Moreover, there are exactly 2 (resp. 1, resp. 0) valid discs linked to $\mathfrak{s}$ if and only if we have $m > 4v(2)$ (resp. $m = 4v(2)$, resp. $m < 4v(2)$), and when $m \geq 4v(2)$ the valid disc(s) containing $\mathfrak{s}$ can be written as $D_{0,2v(2)}$ and $D_{0,m-2v(2)}$ (they coincide when $m = 4v(2)$): this follows from Theorem 6.18 (see also Proposition 6.26). Since we have $\lambda_{\pm}(\mathfrak{s}) = 1$ by Lemma 6.11, by applying Proposition 8.10(b) we can compute that each of the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to D_+ and D_- has abelian rank 0 if $m > 4v(2)$ (the fact that they intersect at 2 nodes follows from Proposition 8.4) and that the component of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to $D_+ = D_- = D_{0,2v(2)}$ is smooth of abelian rank 1 if $m = 4v(2)$. This proves Theorem 9.3(a).

We also remark that, from the formulas for $\mathfrak{t}_{\pm}^{\mathfrak{s}}$ we have derived, one deduces, by Proposition 6.21, that $\mathfrak{t}^{\mathcal{R}}(D_{0,b}) = \min\{b, m-b, 2v(2)\}$ for $b \in [0, m]$, as shown in Figure 1.

1.2. Finding a center and a depth of a valid disc not containing any roots. In the previous subsection, we found all valid discs linked to $\mathfrak{s}$. Moreover, we have seen that, when $m > 4v(2)$, the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ has abelian rank zero, while, if $m = 4v(2)$, it has abelian rank 1, which is entirely contributed by the unique valid disc $D_+ = D_-$ linked to $\mathfrak{s}$. When $m < 4v(2)$, the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ must have abelian rank 1 (as we have $g = 1$ but the toric rank is 0 in the absence of viable clusters by Proposition 8.4), but there is no valid disc linked to $\mathfrak{s}$.

Now, Corollary 8.12 ensures that a valid discs D that is not linked to $\mathfrak{s}$, or that is the unique disc linked to $\mathfrak{s}$, gives a positive contribution to the abelian rank of $(\mathcal{Y}^{\mathrm{rst}})_s$: it follows that there is no such disc when $m > 4v(2)$, and exactly one when $m \leq 4v(2)$. In the latter case, Corollary 7.8 ensures that this unique valid disc D contains all roots α of

$$F(T) = P_1^2(T) - 4P_2(T)P_0(T), \tag{51}$$

where $P_i(T) \in K[T]$ be defined as in the statement of Theorem 9.3(b). Let us therefore assume $m \leq 4v(2)$, let α_1 be any of the roots of F, and let $D = D_{\alpha_1,b_1}$ be the unique valid disc.

Lemma 9.7. In the setting above, we have $v(\alpha_1) = v(\alpha_1 - a_2) = \frac{1}{2}m$ and $v(\alpha_1 - a_3) = 0$.

Proof. This can be proved by directly inspecting the Newton polygon of F. We present a more theoretical proof which separately treats the cases $m = 0$, $0 < m < 4v(2)$, and $m = 4v(2)$.

When $m = 0$, we study the model $\mathcal{Y}_{D'}$ corresponding to the disc $D' = D_{0,0}$. The (normalized) reduction of f has a simple root at $\overline{a_1} = 0, \overline{a_2}, \overline{a_3}$ and ∞. In particular, the trivial decomposition $f = 0^2 + f$ is good at D' by Proposition 4.18, and we have $\mathfrak{t}^{\mathcal{R}}(D') = 0 < 2v(2)$ and $\mu(\mathcal{X}_{D'}, \overline{a_i}) = \mu(\mathcal{X}_{D'}, \infty) = 0$ for $i = 1, 2, 3$ (see §4.6). Now by Proposition 8.13 we have $D \in \mathfrak{D}_P$ for some $P \neq \overline{a_1}, \overline{a_2}, \overline{a_3}, \infty$, which implies, thanks to Proposition 4.6, that $v(\alpha_1 - a_i) = 0$ for all $i = 1, 2, 3$.

When $0 < m < 4v(2)$, we note that $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ is not differentiable at the input $b = \frac{1}{2}m$ with $\mathfrak{t}^{\mathcal{R}}(D_{0,\frac{1}{2}m}) = \frac{1}{2}m < 2v(2)$; from Lemma 8.16 (and Remark 8.17), one deduces that $v(\alpha_1) = \frac{1}{2}m$, which proves the lemma.

When $m = 4v(2)$, given β an element of valuation $\frac{1}{2}m = 2v(2)$ (which is the depth of D), we have that D is the unique valid disc linked to $\mathfrak{s}$, and the conclusion follows from Corollary 7.8. $\square$

From this knowledge of $v(\alpha_1 - a_i)$, by applying Lemma 9.2, one deduces that $\underline{v}_f(D_{\alpha_1,c}) = \sum_{i=1}^{3} v(\alpha_1 - a_i) = m$ for all $c \geq \frac{1}{2}m$. Meanwhile, we know by Proposition 7.5 that there exists a part-square decomposition $f = q^2 + \rho$ that is totally odd with respect to the center α_1 and such that $\rho_{\alpha_1,1}$ has no linear term; this means that $\rho(x) = (x - \alpha_1)^3$, and we thus get $\underline{v}_\rho(D_{\alpha_1,c}) = 3c$ for all $c \in \mathbb{Q}$. We conclude that $\mathfrak{t}^{\mathcal{R}}(D_{\alpha_1,c}) = \min\{3c - m, 2v(2)\}$ for $c \geq \frac{1}{2}m$; hence, the depth b_1 of the valid disc D can now be obtained by solving the equation $3c - m = 2v(2)$ in the variable $c \in [\frac{1}{2}m, +\infty)$ (see Theorem 6.18), which gives $b_1 = \frac{1}{3}(m + 2v(2))$. This proves Theorem 9.3(b).

2. The $g = 2$ case

We now investigate the structure of $(\mathcal{Y}^{\mathrm{rst}})_s$ where Y is a genus-2 hyperelliptic curve; let $Y : y^2 = f(x)$ be the equation of Y, where the polynomial f has degree 5 and satisfies the simplifying assumptions (1), (2) and (3) listed at the beginning of this section; the roots of f will be denoted $a_1 := 0, a_2, \ldots, a_5$. Clearly there may be 0, 1, 2, or 3 even-cardinality clusters among the cluster data associated to f; except for in the last case of 3 even-cardinality clusters, there may be a single cardinality-3 cluster as well.

The below theorem describes our results on the possible structures of $(\mathcal{Y}^{\mathrm{rst}})_s$ depending on various arithmetic conditions, under the assumption that there exists at most one even-cardinality cluster. Actually, the theorem only addresses the case in which the even-cardinality cluster, if it exists, has cardinality 2 and its parent cluster coincides with $\mathcal{R}$, but it may be adapted any other cluster picture having at most one even-cardinality cluster; see Remark 9.10(a) below for more details. To treat the case of more than one even-cardinality cluster, instead see Remark 9.10(b),(c).

Theorem 9.8. Assume that we are in the $g = 2$ situation and retain all of the above assumptions on f. Assume moreover that there are no cardinality-4 clusters and there is at most one cardinality-2 cluster $\mathfrak{s} \subset \mathcal{R}$; if this cluster exists, we denote its relative depth by $m := \delta(\mathfrak{s})$, whereas if there is no even-cardinality cluster, we set $m = 0$. It is clear that $\mathcal{R}$ can contain at most one cardinality-3 cluster $\mathfrak{s}'$; if it exists, we denote its relative depth by $m' := \delta(\mathfrak{s}')$, whereas if there is no cardinality-3 cluster, we set $m' = 0$. We assume that, when both m and m' are > 0, we have $\mathfrak{s} \cap \mathfrak{s}' = \varnothing$.

We label the roots $a_1, \ldots a_5$ of f in such a way that, when $m > 0$, we have $\mathfrak{s} = \{a_1 = 0, a_2\}$, and when $m' > 0$, we have $\mathfrak{s}' = \{a_3, a_4, a_5\}$. Under the assumption that $m > 0$, we write the polynomial $f_-^{\mathfrak{s}}(z)$ (see §6.3.3.2 for its definition) as $1 + M_1 z + M_2 z^2 + M_3 z^3$ and let $w = v(M_1 - 2\sqrt{M_2}) \geq 0$ for some choice of square root of M_2; when $m' > 0$, we have $w = 0$.

Define the polynomial

$$F(T) = (P_1^2(T) - 4P_2(T)P_0(T))^2 - 64P_4(T)P_0^3(T) \in K[T],$$

where $P_i(T)$ is the z^i-coefficient of $f(z+T)$ for $0 \leq i \leq 5$, which we have seen in Remark 7.4 is the polynomial F defined in §7.1. For any root $\alpha \in \bar{K}$ of F, let $f = q^2 + \rho$ be a part-square decomposition that is totally odd with respect to the center α and such that $\rho_{\alpha,1}$ has no linear term (as is guaranteed to exist by Proposition 7.5(a)), and let $\kappa(\alpha)$ be the valuation of the cubic term of $\rho_{\alpha,1}$.

In the language of Proposition 6.26, when $m > 0$, we have $B_{f,\mathfrak{s}} = \max\{4v(2) - w, \frac{8}{3}v(2)\}$. The set of valid discs and the structure of $(\mathcal{Y}^{\mathrm{rst}})_s$ are fully described more precisely as follows. All elements α_i mentioned in parts (b), (c), and (d) below may be chosen to be roots of F, so that in particular $\kappa(\alpha_i)$ is always defined.

(a) Suppose that $m > \frac{8}{3}v(2)$ and $w \geq \frac{4}{3}v(2)$. Then there are exactly 2 valid discs $D_- := D_{0,\frac{2}{3}v(2)}$ and $D_+ := D_{0,m-2v(2)}$, both of which are linked to $\mathfrak{s}$. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of 2 components corresponding to the discs D_- and D_+ which intersect at 2 nodes and have abelian ranks 1 and 0 respectively.

(b) Suppose that $m > 0$ and $4v(2) - m < w < \frac{4}{3}v(2)$. Then there are two valid discs $D_+ := D_{0,m-2v(2)}$ and $D_- := D_{0,2v(2)-w}$ which are linked to $\mathfrak{s}$; their corresponding components of $(\mathcal{Y}^{\mathrm{rst}})_s$ each have abelian rank 0 and intersect each other at 2 points. There is moreover another valid disc D_{α_1,b_1}, which does not contain a root of f; we have $v(\alpha_1 - a_i) = \frac{1}{2}w$ for $i = 1, 2$, $v(\alpha_1 - a_i) = m'$ for $i = 3, 4, 5$ and $b_1 = m' + \frac{1}{3}(w - \kappa(\alpha_1) + 2v(2))$. The corresponding component of $(\mathcal{Y}^{\mathrm{rst}})_s$ has abelian rank 1 and intersects the component corresponding to D_- at 1 node.

(c) Suppose that we have $m > 0$, $w < \frac{1}{2}m$, and $w \leq 4v(2) - m$. Then there are valid discs $D_1 := D_{\alpha_1,b_1}$ and $D_2 := D_{\alpha_2,b_2}$ with $v(\alpha_1 - a_i) = \frac{1}{2}w$ for $i = 1, 2$, $v(\alpha_1 - a_i) = m'$ for $i = 3, 4, 5$, $v(\alpha_2 - a_i) = \frac{1}{2}(m - w)$ for $i = 1, 2$, and $v(\alpha_2 - a_i) = 0$ for $i = 3, 4, 5$, $b_1 = m' + \frac{1}{3}(w - \kappa(\alpha_1) + 2v(2))$, and $b_2 = \frac{1}{3}(m - w - \kappa(\alpha_2) + 2v(2))$. The discs D_i each do not contain a root of f if $w < 4v(2) - m$; when $w = 4v(2) - m$, the disc D_1 does not, but the disc D_2 is the unique valid disc linked to $\mathfrak{s}$ and coincides with $D_{0,m-2v(2)}$. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of 2 components corresponding to the discs D_1 and D_2, each of abelian rank 1, which intersect at 1 node.

(d) Finally, suppose that we have $m = 0$, or $0 < m \leq \min\{2w, \frac{8}{3}v(2)\}$. Then there is a valid disc $D_1 := D_{\alpha_1,b_1}$ with $v(\alpha_1 - a_i) = \frac{1}{4}m$ for $i = 1, 2$, and $v(\alpha_1 - a_i) = 0$ for $i = 3, 4, 5$, and $b_1 \geq v(\alpha_1)$. We have the following subcases.

   (i) Suppose that $\kappa(\alpha_1) < \frac{2}{5}(\frac{1}{2}m + 2v(2))$. Then there is a second valid disc $D_2 := D_{\alpha_2,b_2}$ where α_2 satisfies $v(\alpha_2 - a_i) = \frac{1}{4}m$ for $i = 1, 2$ and $v(\alpha_2 - a_i) = m'$ for $i = 3, 4, 5$, and we have $b_1 = \frac{1}{3}(\frac{1}{2}m - \kappa(\alpha_1) + 2v(2))$ and $b_2 = b_1 + m'$. Neither of the discs D_i contains a root of f. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ consists of 2 components corresponding to the discs D_i, each of abelian rank 1, which intersect at 1 node.

   (ii) Suppose that $\kappa(\alpha_1) \geq \frac{2}{5}(\frac{1}{2}m + 2v(2))$. Then the only valid disc is D_1; it is (the unique valid disc) linked to $\mathfrak{s}$ if $m = \frac{8}{3}v(2)$ but otherwise does not contain a root of f. Its depth is $b_1 = \frac{1}{5}(\frac{1}{2}m + 2v(2))$. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ thus has exactly 1 component, which has abelian rank 2 (so Y attains good reduction in this case).

Remark 9.9. The cases (a), (b), (c) and (d) of Theorem 9.8 (when $m > 0$) correspond to the possible shapes of the function $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ as b ranges in $I(\mathfrak{s}) = [d_-(\mathfrak{s}), d_+(\mathfrak{s})] = [0, m]$, which are described in Figure 3. Note that in cases (a) and (b) we have that $\mathfrak{s}$ is a viable cluster (i.e. $m > B_{f,\mathfrak{s}}$), while in cases (c) and (d) there are no viable clusters.

Remark 9.10. The theorem only treats the situation where there are no cardinality-4 clusters and at most one cardinality-2 cluster which is not contained in a cardinality-3 cluster; here we briefly explain how to treat cases where this hypothesis does not hold.

(a) If we consider a situation where the only even-cardinality cluster $\mathfrak{s}$ has relative depth m and cardinality 4 (instead of 2), then on applying the automorphism i_a as defined in Remark 5.7, where $a \in \mathfrak{s}$ is a root that does not belong to a cardinality-3 cluster, we obtain a cluster picture in which there is a cardinality-2 cluster (and possibly a cardinality-3 cluster disjoint from it), and then using Proposition 6.23 one can derive analogous statements to everything in the above theorem. The rough idea is as follows. For this case, we write $f_+^{\mathfrak{s}}(z) = 1 + M_1 z + M_2 z^2 + M_3 z^3$ and again set $w = v(M_1 - 2\sqrt{M_2})$. Then we again get $B_{f,\mathfrak{s}} = \max\{4v(2) - w, \frac{8}{3}v(2)\}$, and under each of the main hypotheses of parts (a), (b), (c), and (d) we get the same outcome in terms of the structure of $(\mathcal{Y}^{\mathrm{rst}})_s$ (the number of components corresponding to valid discs linked to or not linked to $\mathfrak{s}$, and how they intersect). The valuations of the centers of the discs as well as their depths are given by different formulas, however. In particular, the valid discs $D_\pm$ claimed in parts (a)

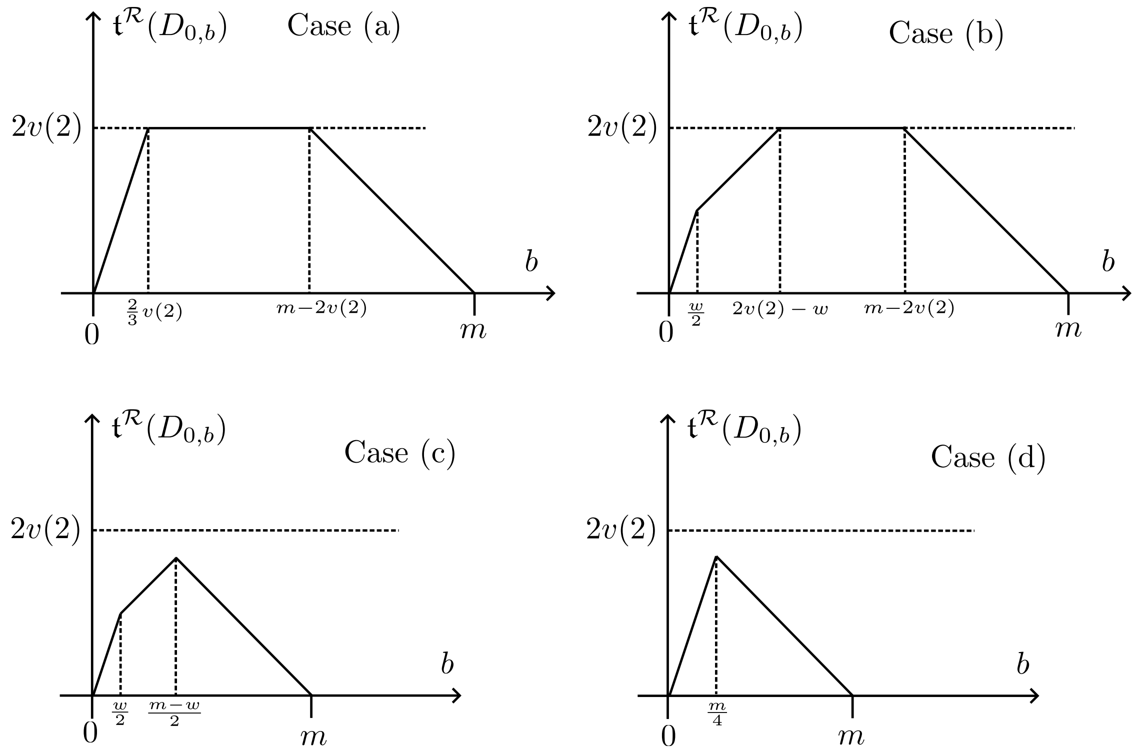

FIGURE 3. The shape of the function $I(\mathfrak{s}) \to [0, 2v(2)]$, $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ in cases (a), (b), (c) and (d) of Theorem 9.8 provided that $m > 0$.

and (b) each have depths $m - b$, where b is the claimed depth in the statement of the theorem: for part (a), we now have valid discs $D_- := D_{a,m-\frac{2}{3}v(2)}$ and $D_+ := D_{a,2v(2)}$ linked to $\mathfrak{s}$, and so on.

Similarly, if we begin with a cluster picture such that there is a cardinality-3 cluster $\mathfrak{s}'$ containing 0, then by applying the automorphism i_a as defined in Remark 5.7 where a is a root in $\mathfrak{s}' \smallsetminus \mathfrak{s}$ (or is any root in $\mathfrak{s}'$ if $m = 0$) and using Proposition 6.23, we obtain a cluster picture in which there is a cardinality-3 cluster which instead does not contain 0.

(b) Suppose that there are exactly 2 even-cardinality clusters $\mathfrak{s}_1$ and $\mathfrak{s}_2$ containing roots a_1 and a_2 respectively. Then by applying appropriate parts of Proposition 6.35(d) combined with Proposition 6.35(a),(b), we get $B_{f,\mathfrak{s}_1} = B_{f,\mathfrak{s}_2} = 4v(2)$. For $i = 1, 2$, the arguments used in §9.1 give us the following statements. If $m_i := \delta(\mathfrak{s}_i) < 4v(2)$ (resp. $\delta(\mathfrak{s}_i) \geq 4v(2)$), then there exist valid discs $D_-^{(1)} := D_{\alpha_i, b_i}$ where α_i is a root of F with $v(\alpha_i - a_i) = d_-(\mathfrak{s}_i) + \frac{1}{2}m_i$ and $b_i = d_-(\mathfrak{s}_i) + \frac{1}{3}(m_i + 2v(2))$ (resp. valid discs $D_+^{(i)} := D_{\mathfrak{s}_i, d_+(\mathfrak{s}_i) - 2v(2)}$ and $D_-^{(i)} := D_{\mathfrak{s}_i, d_-(\mathfrak{s}_i) + 2v(2)}$; these discs coincide if and only if $m_i = 4v(2)$). Moreover, if $m_i \leq 4v(2)$ (resp. if $m_i > 4v(2)$), then the disc $D_-^{(i)}$ contributes a component of the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ of abelian rank 1 (resp. the discs $D_\pm^{(i)}$ each contribute a component of abelian rank 0 and these components meet at 2 nodes), and the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to $D_-^{(1)}$ and $D_-^{(2)}$ intersect at 1 node.

In fact, it is straightforward to compute that, with the quantity w defined as in the theorem, when there are exactly 2 even-cardinality clusters, then we have $w = 0$; the above statements can therefore be proved for each i by applying the below arguments in the proof of part (b) (resp. (c)) of Theorem 9.8 to $\mathfrak{s}_i$ in the case that $m_i > 4v(2)$ (resp. $m_i \leq 4v(2)$) to obtain valid discs $D_\pm^{(i)}$ (resp. the valid disc $D_-^{(i)}$) with the claimed properties.

(c) In the case that there are 3 even-cardinality clusters, the computation of valid discs is in general much less straightforward, but in most cases either Proposition 6.30 or Corollary 6.31 can be applied to entirely determine the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$.

Remark 9.11. Let $\alpha \in \bar{K}$ be a root of F. It is implicit in our proof of Theorem 9.8 that the rational number $\kappa(\alpha)$ is well defined in all contexts of the statement in which its precise value is relevant (more precisely, one can show that it does not depend on the choice of totally odd decomposition with no linear term as long as it is $< 2v(2)$, which is guaranteed to be the case outside of parts (a) and (d)(ii)). We see from the formula for R_3 found in §6.7.7.3 that it can be computed as

$$\kappa(\alpha) = v\Big(P_3(\alpha) - 2\sqrt{P_4(\alpha)}\sqrt{P_2(\alpha) - 2\sqrt{P_4(\alpha)P_0(\alpha)}}\Big) \tag{52}$$

only for particular choices of the square roots in the above formula.

Remark 9.12. We observe the following regarding valuations of roots of the polynomial F.

(a) The polynomial F has degree 16; its leading term has unit coefficient; and its constant term equals $(a_2a_3a_4a_5)^4$, and hence, under the hypotheses of the theorem, it has valuation $4m$.

(b) In light of Corollary 7.8, parts (c) and (d) of the theorem now allow us to deduce the valuations of the roots of the polynomial F. If $m > 0$, $w < \frac{1}{2}m$ and $w \leq 4v(2) - m$, then the statement of Theorem 9.8(c) implies that all roots of F have valuation either $\frac{1}{2}w$ or $\frac{1}{2}(m - w)$; since there are 16 roots whose valuations must add up to $4m$, we get that 8 of the roots have valuation $\frac{1}{2}w$ while the other 8 have valuation $\frac{1}{2}(m - w)$. Similarly, if $m = 0$ or $0 < m \leq \min\{2w, \frac{8}{3}v(2)\}$, then the statement of Theorem 9.8(d) implies that all roots of F have valuation $\frac{1}{4}m$.

Example 9.13. Let Y be the hyperelliptic curve of genus 2 over $\mathbb{Z}_2^{\mathrm{unr}}$ given by

$$y^2 = x(x - 16)(x - 1)(x^2 + x - 1),$$

so that we have a cardinality-2 cluster $\mathfrak{s} = \{0, 16\}$ of relative (and absolute) depth $m = 4v(2)$. It is straightforward to compute that $f_-^{\mathfrak{s}}(z) = 1 - 2z + z^3$ and so we have $w = v(-2 - 2\sqrt{0}) = v(2)$. The hypothesis of Theorem 9.8(b) clearly holds, and so we have valid discs $D_1 := D_- = D_{0,v(2)}$ and $D_2 := D_+ = D_{0,2v(2)}$ which are linked to $\mathfrak{s}$. By applying the computations in §6.7.7.2, we get totally odd decompositions of $f_\pm^{\mathfrak{s},0}$, which induce (as in §6.3.3.2) the decomposition

$$f(x) = [x]^2 + [x^5 - 16x^4 - 2x^3 + 32x^2 - 16x],$$

which according to Proposition 6.42 is good at the discs D_i. Using this decomposition and our knowledge of the depths of the valid discs D_i, following the computations in §4.4 we obtain that the changes in coordinate corresponding to these discs may be written as

$$x = 2x_1 = 4x_2, \quad y = 4y_1 + 2x_1 = 8y_2 + 4x_2. \tag{53}$$

We now get equations for the corresponding models $\mathcal{Y}_1$ and $\mathcal{Y}_2$ as

$$\begin{aligned} \mathcal{Y}_1 &: y_1^2 + x_1y_1 = 2x_1^5 - 2^4x_1^4 - x_1^3 + 2^3x_1^2 - 2x_1 \\ \mathcal{Y}_2 &: y_2^2 + x_2y_2 = 2^4x_2^5 - 2^6x_2^4 - 2x_2^3 + 2^3x_2^2 - x_2 \end{aligned} \tag{54}$$

The special fibers of these models are the $\bar{\mathbb{F}}_2$-curves given by

$$(\mathcal{Y}_1)_s : y_1^2 + x_1y_1 = x_1^3, \qquad y_2^2 + x_2y_2 = x_2. \tag{55}$$

The desingularizations of $(\mathcal{Y}_1)_s$ and $(\mathcal{Y}_2)_s$ are each smooth curves of genus 0 and give rise to 2 of the components of $(\mathcal{Y}^{\mathrm{rst}})_s$. However, these are not all of the components of $(\mathcal{Y}^{\mathrm{rst}})_s$, as Theorem 9.8(b) asserts the existence of another valid disc $D_3 := D_{\alpha_1,b_1}$ for some root α_1 of F with $v(\alpha_1) = \frac{1}{2}v(2)$ and $b_1 = 1 - \frac{1}{3}\kappa(\alpha_1)$. Now through tedious but straightforward calculations, one can show that $v(P_3(\alpha_1)) = v(2)$ and $v(P_4(\alpha_1)) = \frac{1}{2}v(2)$, from which it follows from considering the cubic coefficient appearing in (39) that we have $\kappa(\alpha_1) = v(2)$ and so $b_1 = \frac{2}{3}v(2)$.

For an appropriate part-square decomposition $f = q^2 + \rho$ that is totally odd with respect to the center α_1, the change in coordinates corresponding to D_3 can be written as

$$x = 2^{2/3}x_3 + \alpha_1, \quad y = 2^{3/2}y_3 + q_{\alpha_1,1}(2^{2/3}x_3)y_3.$$

We now get an equation for the model $\mathcal{Y}_3$ corresponding to D_3 as

$$\mathcal{Y}_3 : y_3^2 + 2^{-1/2}q_{\alpha_1,1}(2^{2/3}x_3)y_3 = 2^{-3}\rho_{\alpha_1,1}(2^{2/3}x_3). \tag{56}$$

Note that using Lemma 9.2, we have

(57)

$$v(q_{\alpha_1,1}(2^{2/3}x_3)) = \frac{1}{2}v(f_{\alpha_1,1}(2^{2/3}x_3) - \rho_{\alpha_1,1}(2^{2/3}x_2)) = \frac{1}{2}v(f_{\alpha_1,1}(2^{2/3}x_3)) = v(\alpha_1) = \frac{1}{2}v(2),$$
$$v(\rho_{\alpha_1,1}(2^{2/3}x_3)) = 2v(2) + v(f_{\alpha_1,1}(2^{2/3}x_3)) = 2v(2) + 2v(\alpha_1) = 3v(2).$$

One can now readily verify that the special fiber of $\mathcal{Y}_3$ is the $\bar{\mathbb{F}}_2$-curve given by

$$y_3^2 + c_1y_3 = c_2x_3^3, \tag{58}$$

for some $c_1, c_2 \in k^\times$, and its desingularization is a smooth curve of genus 1 which gives rise to the remaining component of $(\mathcal{Y}^{\mathrm{rst}})_s$. The configuration of the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ is seen in Figure 4.

The rest of this subsection is devoted to proving Theorem 9.8.

2.1. Finding valid discs containing an even-cardinality cluster. Suppose that, in the settting of Theorem 9.8, we have $m > 0$, i.e. that we have a unique even-cardinality cluster $\mathfrak{s} := \{0, a_2\}$ of relative depth m; our goal for the moment is to find all valid discs which are linked to $\mathfrak{s}$. We adopt the notation and constructions of §6.3 and get the polynomials $f_+^{\mathfrak{s}}(z) = 1 - z$ and

$$f_-^{\mathfrak{s}}(z) = (1 - a_3^{-1}z)(1 - a_4^{-1}z)(1 - a_5^{-1}z) = 1 + M_1z + M_2z^2 + M_3z^3. \tag{59}$$

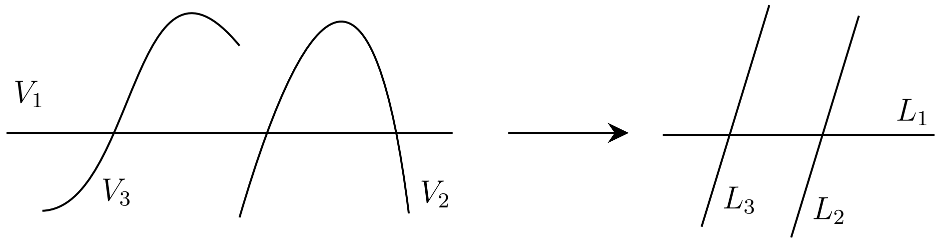

FIGURE 4. The special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$, shown on the left, mapping to $(\mathcal{X}^{(\mathrm{rst})})_s$; each component V_i of $(\mathcal{Y}^{\mathrm{rst}})_s$ maps to each component $L_i := (\mathcal{X}_{D_i})_s$ of $(\mathcal{X}^{(\mathrm{rst})})_s$.

Just as in the situation of §9.1.1.1, we have $\mathfrak{t}_+^{\mathfrak{s}}(b) = \min\{b, 2v(2)\}$ and $b_0(\mathfrak{t}_+^{\mathfrak{s}}) = 2v(2)$. Now applying the computations in §6.7.7.2, we have a totally odd part-square decomposition $f_-^{\mathfrak{s},0} = [q_-]^2 + \rho_-$ where (for some choice of square roots of M_2) we have

$$\rho_-(z) = (M_1 - 2\sqrt{M_2})z + M_3 z^3.$$

It is immediate to see that we have $v(M_3) = 0$, and the formula $\mathfrak{t}_-^{\mathfrak{s}}(b) = \min\{3b, b + w, 2v(2)\}$ follows, from which we get $b_0(\mathfrak{t}_-^{\mathfrak{s}}) = \max\{\frac{2}{3}v(2), 2v(2) - w\}$. Now, using the formula $B_{f,\mathfrak{s}} = b_0(\mathfrak{t}_+^{\mathfrak{s}}) + b_0(\mathfrak{t}_-^{\mathfrak{s}})$ from Proposition 6.26, we get $B_{f,\mathfrak{s}} = \max\{\frac{8}{3}v(2), 4v(2) - w\}$. If we assume that $\mathfrak{s}$ is viable (i.e., if we are in the setting $m > \max\{\frac{8}{3}v(2), 4v(2) - w\}$ treated by Theorem 9.8(a),(b)), the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to $D_\pm := D_{0,b_\pm(\mathfrak{s})}$ intersect at 2 nodes (see Proposition 8.4). Moreover, it follows from Lemma 6.11 that $\lambda_+(\mathfrak{s}) = 1$, and it is easily checked from the valuations of the coefficients of ρ_- that we have $\lambda_-(\mathfrak{s}) = 3$ if we moreover have $w \geq \frac{4}{3}v(2)$ (i.e. case (a)), whereas $\lambda_-(\mathfrak{s}) = 1$ if $w < \frac{4}{3}v(2)$ (i.e. case (b)). Now applying Proposition 8.10(b)(i),(ii) shows us that in case (a) of Theorem 9.8, the components of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to D_+ and D_- have abelian ranks 0 and 1 respectively if $m > \frac{8}{3}v(2)$.

In particular, we have proved the formula for $B_{f,\mathfrak{s}}$ at the start of the statement of Theorem 9.8, as well as part (a) of the theorem and part of the statement of part (b). We also note for below use that from the formulas $\mathfrak{t}_+^{\mathfrak{s}}(b) = \min\{b, 2v(2)\}$ and $\mathfrak{t}_-^{\mathfrak{s}}(b) = \min\{3b, b + w, 2v(2)\}$, by Proposition 6.21, we get the formula

$$\mathfrak{t}^{\mathcal{R}}(D_{0,b}) = \min\{3b, b + w, -b + m, 2v(2)\} \qquad \text{for } b \in [d_-(\mathfrak{s}), d_+(\mathfrak{s})] = [0, m], \tag{60}$$

as displayed in Figure 3.

2.2. Finding a center and a depth of a valid disc not containing any roots. We retain all of the above notation and assumptions, except that we now allow the possibility that $m = 0$ (so that there is no even-cardinality cluster), and we set out to find and chararcterize the valid discs D associated to Y which either are linked to no cluster or are the unique valid disc linked to $\mathfrak{s}$.

Below we will need a lemma to treat situations where $w = 0$ (which is possible only under the hypotheses of Theorem 9.8(b),(c)).

Lemma 9.14. With the notation and hypotheses of *Theorem* 9.8, suppose that we have $m > 0$ and $w = 0$. Then there is a valid disc D containing no roots of f and which, for all $\alpha \in D$, satisfies

$$v(\alpha) = v(\alpha - a_2) = 0, \qquad v(\alpha - a_3) = v(\alpha - a_4) = v(\alpha - a_5) = m'. \tag{61}$$

Moreover, D contributes a component of abelian rank 1 to the special fiber $(\mathcal{Y}^{\mathrm{rst}})_s$ that meets the rest of $(\mathcal{Y}^{\mathrm{rst}})_s$ at 1 node.

Proof. Let $D' := D_{0,0}$. The hypothesis $w = 0$, by definition of w, is equivalent to $v(M_1) = 0$, and one checks straightforwardly from formulas that it implies that f has unit cubic coefficient. Moreover, the polynomial f has unit quintic coefficient (because it is monic), while the presence of the cardinality-2 cluster $\mathfrak{s}$ implies that the linear term of f has positive valuation. Hence, the decomposition $f = 0^2 + f$ is good at D'. Moreover, when $m' = 0$, by looking at the roots of $\bar{f}'$ we see that the inseparable curve $(\mathcal{Y}_{D'})_s \to (\mathcal{X}_{D'})_s$ is singular exactly over $\overline{x} = 0$ and over a second point P to which none of the elements of $\mathcal{R} \cup \{\infty\}$ reduce, with $\mu(\mathcal{X}_{D'}, 0) = \mu(\mathcal{X}_{D'}, P) = 2$ (see §4.6). Now applying Proposition 8.14 (combined with Proposition 4.6), we get the desired statement when $m' = 0$. When we have $m' > 0$, we instead let $D' = D_{a_i,m'}$ for $i = 3, 4, 5$, and letting γ' be an element of valuation m', one easily sees that any normalized reduction of $f_{a_i,\gamma'}$ has no quintic term but does have nonzero linear and cubic terms. It follows that $(\mathcal{Y}_{D'})_s \to (\mathcal{X}_{D'})_s$ is singular exactly over ∞ and over a second point P to which none of the elements of $\mathcal{R} \cup \{\infty\}$ reduce, with $\mu(\mathcal{X}_{D'}, 0) = \mu(\mathcal{X}_{D'}, P) = 2$. Again, the desired statement follows via Proposition 8.14. $\square$

In the case treated by Theorem 9.8(a), where $m > \frac{8}{3}v(2)$ and $w \geq \frac{4}{3}v(2)$, we have seen that there are 2 valid discs linked to $\mathfrak{s}$: they have already been found and determined to contribute 1 to the abelian rank and 1 to the toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$; since we have $g = 2$ (so that the sum of the ranks must equal 2; see §2.1.1.7) and valid discs not linked to any cluster correspond to components of positive abelian rank by Corollary 8.12, it is clear that there is no valid disc which does not contain a root of f or which is the unique one linked to $\mathfrak{s}$. We therefore assume that the hypothesis of Theorem 9.8(a) does not hold.

Suppose that we have $m > 0$ and $w < \min\{\frac{1}{2}m, \frac{4}{3}v(2)\}$ (as is true for the cases treated by Theorem 9.8(b),(c)); we will show that there is a valid disc D_{α_1,b_1} containing no root of f such that α_1 satisfies

$$v(\alpha_1) = v(\alpha_1 - a_2) = \frac{1}{2}w, \text{ and } v(\alpha_1 - a_3) = v(\alpha_1 - a_4) = v(\alpha_1 - a_5) = m' \tag{62}$$

and which contributes a component of abelian rank 1 to $(\mathcal{Y}^{\mathrm{rst}})_s$. If $w = 0$, this follows from Lemma 9.14. We therefore assume for the rest of this paragraph that $w > 0$. Then we have that $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,c})$ is not differentiable at the input $c = \frac{1}{2}w$ with left and right derivatives equal to 3 and 1 respectively and that $\mathfrak{t}^{\mathcal{R}}(D_{0,\frac{1}{2}w}) = \frac{3}{2}w < 2v(2)$. Therefore, by Lemma 8.16 (and Remark 8.17), there is a center $\alpha_1 \in \bar{K}$ such that $v(\alpha_1) = \frac{1}{2}w$ and D_{α_1,b_1} is a valid disc which is not linked to any cluster for some $b_1 > \frac{1}{2}w$. In fact, we know from the left and right derivatives and applying Proposition 8.14(b) to the disc $D_{0,\frac{1}{2}w}$ that the abelian rank of the corresponding component of $(\mathcal{Y}^{\mathrm{rst}})_s$ must be $\frac{1}{2}(3 - 1) = 1$. When we also have $w > 4v(2) - m$ (so that we are in the case treated by Theorem 9.8(b)), we have already found that $\mathfrak{s}$ is viable and contributes 1 to the toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$ and so we have found all of the valid discs as the ranks must add up to $g = 2$ (see §2.1.1.7).

In the case treated by Theorem 9.8(c) where we moreover have $w \leq 4v(2) - m$ and $w < \frac{1}{2}m$, the function $c \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,c})$ is also not differentiable at the input $c = \frac{1}{2}(m - w)$ with left and right derivative equal to 1 and -1 respectively, and we have $\mathfrak{t}^{\mathcal{R}}(D_{0,\frac{1}{2}(m-w)}) = \frac{1}{2}(m + w) \leq 2v(2)$. Therefore, by applying Lemma 8.16 and Remark 8.17 when $w < 4v(2) - m$, and by observing that $b_-(\mathfrak{s}) = b_+(\mathfrak{s}) = \frac{1}{2}(m - w)$ when $w = 4v(2) - m$, we have that there is a center α_2 satisfying

$$v(\alpha_2) = v(\alpha_2 - a_2) = \frac{1}{2}(m - w) > 0, \text{ and } v(\alpha_2 - a_3) = v(\alpha_2 - a_4) = v(\alpha_2 - a_5) = 0, \tag{63}$$

and such that D_{α_2,b_2} is a valid disc which is not linked to any cluster (resp. is the unique valid disc linked to $\mathfrak{s}$) if $w < 4v(2) - m$ (resp. $w = 4v(2) - m$), for some $b_2 \geq \frac{1}{2}(m - w)$. By knowing the left and right derivatives and a similar application of Proposition 8.14(b) to the disc $D_{0,\frac{1}{2}(m-w)}$ (resp. using Proposition 8.10(b)(iii)), the abelian rank of the corresponding component of $(\mathcal{Y}^{\mathrm{rst}})_s$ must be $\frac{1}{2}(1 - (-1)) = 1$. Since the 2 valid discs we have found each contribute 1 to the abelian rank of $(\mathcal{Y}^{\mathrm{rst}})_s$, we have again found all of the valid discs as $g = 2$ (see §2.1.1.7).

Let us now address the case treated by Theorem 9.8(d), in which we instead have $m = 0$ or $0 < m \leq \min\{2w, \frac{8}{3}v(2)\}$ and that the function $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{0,b})$ is not differentiable at the input $b = \frac{1}{4}m$, with $\mathfrak{t}^{\mathcal{R}}(D_{0,\frac{1}{4}m}) = \frac{3}{4}m \leq 2v(2)$. We claim that, in this case, there exist two possibly coinciding valid discs D_{α_1,b_1} and D_{α_2,b_2} such that α_1 and α_2 satisfy the conditions

$$v(\alpha_1) = v(\alpha_1 - a_2) = \frac{1}{4}m, \text{ and } v(\alpha_1 - a_3) = v(\alpha_1 - a_4) = v(\alpha_1 - a_5) = 0, \tag{64}$$

$$v(\alpha_2) = v(\alpha_2 - a_2) = \frac{1}{4}m, \text{ and } v(\alpha_2 - a_3) = v(\alpha_2 - a_4) = v(\alpha_2 - a_5) = m' \tag{65}$$

respectively, and such that they are linked to no cluster (when $m < \frac{8}{3}v(2)$), or they both coincide with the unique disc linked to $\mathfrak{s}$ (when $m = \frac{8}{3}v(2)$). When $m = \frac{8}{3}v(2)$, it is straightforward to see that the unique valid disc $D_{0,\frac{2}{3}v(2)} = D_{\alpha_1,b_1} = D_{\alpha_2,b_2}$ linked to $\mathfrak{s}$ contributes a component of abelian rank 2 (see Proposition 8.10(b)(iii)), which proves the claim directly above. In the $m < \frac{8}{3}v(2)$ case, no valid disc is linked to $\mathfrak{s}$, and hence all valid discs contain no roots of f. Since the genus is $g = 2$ and the toric rank of $(\mathcal{Y}^{\mathrm{rst}})_s$ is 0 by Theorem 8.1, in light of Corollary 8.12 we must have either two distinct valid discs each contributing a component of abelian rank 1 or a unique valid disc contributing a component of abelian rank 2. The fact that these valid discs have the form prescribed by the claim directly above can be easily proved by studying the singularities of $(\mathcal{Y}_{D'})_s \to (\mathcal{X}_{D'})_s$ for the disc $D' = D_{0,\frac{m}{4}}$ (and also for the disc $D' = D_{0,m'}$ when $m' > 0$) and then exploiting Proposition 8.13.

We now set out to find a formula for b_i for the valid discs D_{α_i,b_i} that we have found in all cases discussed above, where $i = 1, 2$ and the centers α_i satisfy (62)-(63) or (64)-(65). We also want to know whether, in the case that $m = 0$ or $0 < m < \min\{2w, \frac{8}{3}v(2)\}$, the valid discs D_{α_1,b_1} and D_{α_2,b_2} give rise to two distinct components of $(\mathcal{Y}^{\mathrm{rst}})_s$ of abelian rank 1 or they coincide and provide a unique component of abelian rank 2.

Let us first remark that α_i can always be chosen to be a root of F thanks to Corollary 7.8. Now, using Lemma 9.2 and the definition of $\kappa(\alpha)$, letting $\nu = 2v(\alpha_i) + 3v(\alpha_i - a_3)$, we compute the formula

$$\mathfrak{t}^{\mathcal{R}}(D_{\alpha_i,b}) = \min\{\underline{v}_\rho(D_{\alpha_i,b}) - \underline{v}_f(D_{\alpha_i,b}), 2v(2)\} = \min\{3b + \kappa(\alpha_i) - \nu, 5b - \nu, 2v(2)\}. \tag{66}$$

We know from Theorem 6.18(b) that the depth b_i is the first input at which $b \mapsto \mathfrak{t}^{\mathcal{R}}(D_{\alpha_i,b})$ attains $2v(2)$. Meanwhile, it can be calculated using Proposition 8.10(a) and the formula in (66) that the component of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to D_{α_i,b_i} has abelian rank 1 (resp. 2) if we have $3b_i + \kappa(\alpha_i) - \nu < 5b_i - \nu$ (resp. $3b_i + \kappa(\alpha_i) - \nu \geq 5b_i - \nu$), or equivalently, if we have $\kappa(\alpha_i) < 2b_i$ (resp. $\kappa(\alpha_i) \geq 2b_i$). We therefore have have the following two cases:

(1) the abelian rank of the component of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to D_{α_i,b_i} equals 1; we then have

$$\begin{aligned} \mathfrak{t}^{\mathcal{R}}(D_{\alpha_i,b_i}) = 3b_i + \kappa(\alpha_i) - \nu = 2v(2) &\implies b_i = \frac{1}{3}(\nu - \kappa(\alpha_i) + 2v(2)), \\ \text{and} \quad \kappa(\alpha_i) < 2b_i &\implies \kappa(\alpha_i) < \frac{2}{5}(\nu + 2v(2)); \qquad \text{or} \end{aligned} \tag{67}$$

(2) the abelian rank of the component of $(\mathcal{Y}^{\mathrm{rst}})_s$ corresponding to D_{α_i,b_i} equals 2; we then have

$$\begin{aligned} \mathfrak{t}^{\mathcal{R}}(D_{\alpha_i,b_i}) = 5b_i - \nu = 2v(2) &\implies b_i = \frac{1}{5}(\nu + 2v(2)), \\ \text{and} \quad \kappa(\alpha_i) \geq 2b_i &\implies \kappa(\alpha_i) \geq \frac{2}{5}(\nu + 2v(2)). \end{aligned} \tag{68}$$

Now, on substituting the formulas for $\nu = 2v(\alpha_1) + 3v(\alpha_1 - a_3)$ which we found above on a case-by-case basis, we get the following outcomes:

(1) if we have $w < \min\{\frac{1}{2}m, \frac{4}{3}v(2)\}$, then the valid disc D_{α_1,b_1} that we have found (corresponding to a component of abelian rank 1), with α_1 satisfying (62), has depth $b_1 = m' + \frac{1}{3}(w - \kappa(\alpha_1) + 2v(2))$;

(2) if moreover we have $w < \min\{\frac{1}{2}m, 4v(2) - m\}$, then the second valid disc D_{α_2,b_2} that we have found (also corresponding to a component of abelian rank 1), with α_2 satisfying (63), has depth $b_2 = \frac{1}{3}(m - w - \kappa(\alpha_2) + 2v(2))$; and

(3) if instead we have $m = 0$ or $0 < m < \min\{2w, \frac{8}{3}v(2)\}$, then the valid discs D_{α_i,b_i} for $i = 1, 2$ that we have found, with α_i satisfying (64)-(65), may be distinct and contribute each a component of abelian rank 1, or they may coincide and give a unique component of abelian rank 2. The first case occurs if we have $\kappa(\alpha_1) < \frac{2}{5}(\frac{1}{2}m + 2v(2))$; in this case, we get $b_1 = \frac{1}{3}(\frac{1}{2}m - \kappa(\alpha_1) + 2v(2))$ and $b_2 = m' + b_1$. The second case (which can only happen if $m' = 0$) occurs when $\kappa(\alpha_1) \geq \frac{2}{5}(\frac{1}{2}m + 2v(2))$, and we get $b_1 = b_2 = \frac{1}{5}(\frac{1}{2}m + 2v(2))$.

This finishes the proof of Theorem 9.8.

Bibliography

[1] Michael Artin and Gayn Winters. "Degenerate fibres and stable reduction of curves". In: *Topology* 10.4 (1971), pp. 373–383.

[2] Siegfried Bosch, Werner Lütkebohmert, and Michel Raynaud. *Néron models*. Vol. 21. Springer Science & Business Media, 2012.

[3] Irene I Bouw and Stefan Wewers. "Computing L-functions and semistable reduction of superelliptic curves". In: *Glasgow Mathematical Journal* 59.1 (2017), pp. 77–108.

[4] Irene I. Bouw and Stefan Wewers. "Semistable reduction of curves and computation of bad Euler factors of L-functions". In: ICERM course notes (2015).

[5] Robert Coleman. "Computing stable reductions". In: *Séminaire de Théorie des Nombres, Paris 1985–86*. Springer, 1987, pp. 1–18.

[6] Pierre Deligne and David Mumford. "The irreducibility of the space of curves of given genus". In: *Publications Mathématiques de l'IHES* 36 (1969), pp. 75–109.

[7] Tim Dokchitser et al. "Arithmetic of hyperelliptic curves over local fields". In: *Mathematische Annalen* (385 2023), pp. 1213–1322.

[8] Vladimir Dokchitser and Adam Morgan. "A note on hyperelliptic curves with ordinary reduction over 2-adic fields". In: *arXiv preprint arXiv:2203.11254* (2022).

[9] Tim Gehrunger and Richard Pink. "Reduction of Hyperelliptic Curves in Characteristic $\neq 2$". In: *arXiv preprint arXiv:2112.05550* (2021).

[10] Claus Lehr, Michel Matignon, et al. "Wild monodromy and automorphisms of curves". In: *Duke Mathematical Journal* 135.3 (2006), pp. 569–586.

[11] Qing Liu. *Algebraic geometry and arithmetic curves*. Vol. 6. Oxford University Press on Demand, 2002.

[12] Qing Liu and Dino Lorenzini. "Models of Curves and Finite Covers". In: *Compositio Mathematica* 118.1 (1999), pp. 61–102.

[13] Michel Matignon. "Vers un algorithme pour la réduction stable des revêtements p-cycliques de la droite projective sur un corps p-adique". In: *Mathematische Annalen* 325.2 (2003), pp. 323–354.

[14] Jean-Pierre Serre. *Abelian ℓ-adic representations and elliptic curves*. Addison-Wesley, Advanced Book Program (Redwood City, Calif.), 1989.

[15] Igor Shafarevich. *Basic algebraic geometry: Varieties in projective space*. Springer, 1994.

[16] Joseph H. Silverman. "*The arithmetic of elliptic curves*". In: *Graduate Texts in Mathematics* 106 (2009).

[17] Jeffrey Yelton. "Semistable models of elliptic curves over residue characteristic 2". In: *Canadian Mathematical Bulletin* 64.1 (2021), pp. 154–162.

Part 2

Theta operators

CHAPTER 10

Introduction

1. Theta operators

In 1977, Katz constructed a weight-raising differential operator on the space of p-adic modular forms, known as the ϑ operator. At the level of q-expansions, the operator ϑ can be described as $q\frac{d}{dq}$, meaning that $\vartheta(\sum_n a_n q^n) = \sum_n n a_n q^n$.

The construction that Katz presents in [7] goes as follows: if $X_1(N) \to \mathrm{Spec}(\mathbb{Z}[1/N])$ is the compactified modular curve of level $\Gamma_1(N)$ for some $N \geq 4$, and $\mathfrak{X}_1(N) \to \mathrm{Spf}(\mathbb{Z}_p)$ denotes its p-adic completion at a prime $p \nmid N$, one has that, over the ordinary locus $\mathfrak{X}_1^{\mathrm{ord}}(N) \subset \mathfrak{X}_1(N)$, the Hodge filtration $0 \to \omega \to H^1_{\mathrm{dR}} \to \omega^\vee \to 0$ splits canonically, via the so-called unit-root splitting, and this allows to derive a differential operator $\vartheta : \omega^k \to \omega^{k+2}$ (for any $k \in \mathbb{Z}$) from the Gauss-Manin connection on H^1_{dR}. This operator has some interesting commutation properties with Hecke operators T_ℓ and U_ℓ at each prime ℓ:

$$T_\ell \vartheta = \ell T_\ell \vartheta \quad \text{for } \ell \nmid N;$$
$$U_\ell \vartheta = \ell U_\ell \vartheta \quad \text{for } \ell \mid N.$$

As a consequence, applying the theta operator to a modular form corresponds to twisting by a cyclotomic character its associated Galois representation, as the following result formalizes (see, for example, the introduction of [4] for a more detailed discussion).

Theorem. Let f be a normalized cuspidal eigenform of weight $k \in \mathbb{Z}/(p-1)\mathbb{Z}$, level $\Gamma_1(N)$ and nebentypus $\varepsilon : (\mathbb{Z}/N\mathbb{Z})^\times \to E^\times$, with coefficients in some finite extension $E/\mathbb{F}_p$. Let $\rho_f : G_\mathbb{Q} \to \mathrm{GL}_2(E)$ be the continuous semisimple Galois representation attached to it. If $g := \vartheta(f)$, then g is still a normalized cuspidal eigenform with coefficients in E for the group $\Gamma_1(N)$, its weight is $k+2$, and $\rho_g \cong \rho_f \otimes \omega$, being $\omega : G_\mathbb{Q} \to \mathrm{GL}_1(E)$ the cyclotomic character.

Recent progress has been done in reproducing Katz's construction for more general Shimura varieties, and in the consequent study of the resulting differential operators, both in the mod p and in the p-adic setting. In 2012, E. Eishen published her work [2], which discusses how to apply Katz's construction to Shimura varieties attached to unitary groups of the form $U(n,n)$. A generalization of the result appeared in the 2018 paper [4] by Eishen, Fintzen, Mantovan and Varma, which addresses more general Shimura varieties of type A and C having non-empty ordinary locus.

When the ordinary locus is empty, one can still adopt a definition similar to Katz's one on the μ-ordinary locus, after suitably replacing the unit-root splitting, which is no longer available, with a subtler canonical splitting of H^1_{dR}, associated to the slope filtration of the universal μ-ordinary p-divisible group. This approach was explored by De Shalit and Goren in [1], and by Eischen and Mantovan in [3].

A quite different perspective on ϑ operators was proposed in the paper [6] by S. Howe, appeared in 2020. The author restricts its attention to the elliptic case, and starts by considering the big Igusa tower on the p-adic modular curve $\mathfrak{X}_1(N) \to \mathrm{Spf}(\mathbb{Z}_p)$, i.e. the moduli space $\mathfrak{IG} \to \mathfrak{X}_1(N)$ classifying the trivialization $A[p^\infty] \cong \mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$ of the universal p-divisible group $A[p^\infty]$. On $\mathfrak{IG}$, the group of self-p-quasi-isogenies of $\mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$ acts, which yields a residual infinitesimal action of the p-divisible group μ_{p^∞} on Katz's space of modular forms. By taking the derivative of this action, one recovers Katz's ϑ operator.

Howe's approach has several advantages. Its geometric nature, together with its independence from q-expansion computations, makes is well-suited for being abstracted and applied to more general Shimura varieties of PEL type. Moreover, the fact that the theta operator is obtained as the differential of a μ_{p^∞}-action ensures that it can be iterated p-adically, without any need of proving this via explicit congruences between powers of ϑ.

In the present work, Howe's point of view is developed in two different directions. First, we adopt a variant of Howe's approach to construct ϑ operators on the μ-ordinary locus of PEL-type Shimura varieties of type A and C. The second construction we perform is only presented in the specific case of the Siegel threefold, where we build a ϑ operator on a sheaf of p-adic modular forms over a formal open that is larger than the ordinary locus (namely, the p-rank ≥ 1 locus).

2. Theta operators on the μ-ordinary locus

Let us consider an integral PEL datum $\mathcal{D}$ (of type A or C) with good reduction at p and a level K (hyperspecial at p, and neat). Denoting by E its reflex field, we consider the integral Shimura variety $Y \to \mathrm{Spec}(\mathcal{O}_{E,(p)})$, and its formal completion $\mathfrak{Y} \to \mathrm{Spf}(\mathcal{O}_{E,p})$. The Shimura variety Y can be characterized as the representative of a moduli problem for $\mathcal{D}$-enriched abelian schemes, and consequently carries a universal abelian scheme $A \to Y$. The μ-ordinary locus $\mathfrak{Y}^{\mu-\mathrm{ord}} \subset \mathfrak{Y}$ is an open and dense formal subscheme of $\mathfrak{Y}$. Over this locus, the universal p-divisible group $A[p^\infty]$ admits a slope filtration, whose extreme graded pieces are the multiplicative one $A[p^\infty]^{\lambda=0}$ and the étale one $A[p^\infty]^{\lambda=1}$: the former is étale-locally isomorphic to $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$ for some appropriate $\mathbb{Z}_p$-lattice $\mathfrak{L}$ with $\mathcal{O}_B$-action (where $\mathcal{O}_B$ is the order providing the endomorphism structure to the enriched abelian scheme A), while the latter is isomorphic to its dual, i.e. $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$. The remaining graded pieces $A[p^\infty]^\lambda$, for $0 < \lambda < 1$, are also étale-locally isomorphic each to some connected-connected fixed p-divisible groups $\mathfrak{X}^\lambda$ defined over $\mathrm{Spec}(\mathcal{O}_{E,(p)})$.

Letting $\mathfrak{M}^{\mathrm{gr}}$ be the group of automorphisms of $\mathbb{X}$ respecting its slope decomposition and its polarization, we can factor $\mathfrak{M}^{\mathrm{gr}}$ as $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \times \mathfrak{M}'$, where $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$ acts on the multiplicative and étale pieces, while $\mathfrak{M}'$ is the automorphism group of the connected-connected part. Given R some p-adically complete $\mathcal{O}_E$-algebra, for each representation $\rho : \mathfrak{M}^{\mathrm{gr}} \to \mathrm{GL}_r(R)$ of the p-adic group $\mathfrak{M}^{\mathrm{gr}}$, one can form a sheaf of *p-adic modular forms of weight ρ over the μ-ordinary locus* (as discussed in [3]). The following result (which is Theorem 12.4) describes the ϑ operators we construct on these sheaves of p-adic modular forms.

Theorem (A). Let $\lambda : \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \to \mathrm{GL}(V)$ and $\rho : \mathfrak{M}^{\mathrm{gr}} \to \mathrm{GL}(W)$ be continuous representations of $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$ and $\mathfrak{M}^{\mathrm{gr}} = \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \times \mathfrak{M}'$ respectively, being V and W some R-vector bundles, for R some p-adically complete $\mathcal{O}_E$-algebra. Let $f : \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}) \to V$ be an $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$-equivariant continuous function. Then, there exists an operator $\vartheta^f : \omega^\rho \to \omega^{\rho\otimes\lambda}$ over $\mathfrak{Y}^{\mu-\mathrm{ord}} \hat{\otimes} R$.

The construction of the operator follows a variant of Howe's approach: we construct a small Igusa tower $\mathfrak{T}^{\mathrm{gr}} \to \mathfrak{Y}^{\mu-\mathrm{ord}}$, on which the isoclinic pieces of the universal p-divisible group are trivialized. Then, we construct an action of the group $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$ on $\mathfrak{T}^{\mathrm{gr}}$, that essentially consists in applying infinitesimal deformations to the universal p-divisible group, and we study its interaction with the $\mathfrak{M}^{\mathrm{gr}}$-action on $\mathfrak{T}^{\mathrm{gr}}$. Finally, we take the differential of the $(\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty})$-action, and we obtain ϑ.

3. Theta operators beyond the μ-ordinary locus: the Siegel threefold case

The Siegel threefold can be seen as the moduli of principally polarized abelian surfaces (with some prime-to-p level structure, that we will always consider fixed, and for which we presume that the neatness assumption is satisfied). Let X be a toroidal compactification of the Siegel threefold, $A \to X$ be the semiabelian surface it carries, and ω the sheaf of invariant differentials of A. For each dominant weight $(k_1, k_2) \in \mathbb{Z}^2$, with $k_1 \geq k_2$, we denote by $\omega^{(k_1,k_2)}$ the vector bundle of modular forms of weight (k_1, k_2) on X. We denote by $\mathfrak{X} \to \mathrm{Spf}(\mathbb{Z}_p)$ the p-adic completion of X.

In his 2020 work [11], devising the foundations of higher Hida and Coleman theories in the case of GSp_4, Pilloni is concerned with defining and studying spaces of p-adic modular forms defined over the p-rank ≥ 1 locus, $\mathfrak{X}^{\geq 1}$. This is larger that the ordinary locus (which can be described as the p-rank 2 locus), it is still an open formal subscheme of $\mathfrak{X}$, and it is not affine: it is, instead, the union of two open formal affine subschemes. The non-affineness of $\mathfrak{X}^{\geq 1}$ makes it interesting, for example, for p-adically deforming those Hecke eigenclasses that lie in the H^1 groups of modular sheaves – which would vanish if restricted to the ordinary locus.

Pilloni's paper [11] proposes a definition of p-adic modular over the locus $\mathfrak{X}^{\geq 1}$ which goes as follows. The first step is constructing an affine pro-étale cover $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)} \to \mathfrak{X}^{\geq 1}$ classifying multiplicative height-1 p-divisible subgroups $H \hookrightarrow A$. Over this deep Klingen cover, Pilloni defines a (small) Igusa variety $\mathfrak{Ig}$ classifying the trivializations $H \cong \mu_{p^\infty}$, and he exploits it to define, for each p-adic character k of $\mathbb{Z}_p^\times$, a sheaf of p-adic modular forms $\mathfrak{F}^k$ as the sheaf of functions on $\mathfrak{Ig}$ on which $\mathbb{Z}_p^\times$ acts via the character k.

Given k_1 a p-adic character of $\mathbb{Z}_p^\times = \mathrm{Aut}(\mu_{p^\infty})$ and $k_2 \in \mathbb{Z}$, we will adopt the notation $\mathfrak{F}^{(k_1,k_2)}$ to denote the sheaf $\mathfrak{F}^{k_1-k_2} \otimes \det(\omega)^{k_2}$: this can be thought of as a p-adic interpolation, over $\mathfrak{X}^{\geq 1}$, of the classical sheaves of modular forms $\{\omega^{(k_1,k_2)} : (k_1, k_2) \in \mathbb{Z}^2, k_1 \geq k_2\}$.

In this work, we show that a generalization of Howe's techniques allows to obtain a differential operator $\mathfrak{F}^{(k_1,k_2)}$ over $\mathfrak{X}^{\geq 1}$. More precisely, we have the following result, which is Theorem 13.4.

Theorem (B). For any pair (k_1, k_2) with $k_1 : \mathbb{Z}_p^\times \to R$ a continuous character (for R some p-adically complete $\mathbb{Z}_p$-algebra) and $k_2 \in \mathbb{Z}$, each multiplicative continuous function $f : \mathbb{Z}_p \to R$ induces an R-linear operator $\vartheta^f : \mathfrak{F}^{(k_1,k_2)} \to \mathfrak{F}^{(k_1+2f,k_2)}$ over $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)} \hat{\otimes} R$ minus the boundary. In particular, for $n \in \mathbb{N}$, the function $\mathbb{Z}_p \to \mathbb{Z}_p$, $x \mapsto x^n$ induces a n-th order R-linear differential operator $\vartheta^n : \mathfrak{F}^{(k_1,k_2)} \to \mathfrak{F}^{(k_1+2n,k_2)}$.

Our construction is analogous to the one we have presented in the previous section for the μ-ordinary locus of PEL-type Shimura varieties. First, we consider an appropriate (small) Igusa covering of $\mathfrak{X}_{\mathrm{Kli}}(p^\infty)$, that we denote $\mathfrak{T}^{\mathrm{gr}}$ and is a $(\mathbb{Z}_p^\times \times \mathrm{GL}_1)$-torsor over $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$. We then exhibit a study a μ_{p^∞}-action on $\mathfrak{T}^{\mathrm{gr}}$, that acts by infinitesimally deforming the p-divisible group $A[p^\infty]$ together with all the additional structure it carries. We finally take take the differential of this action, to obtain the ϑ operator.

This newly constructed ϑ operator has suitable commutation properties with Hecke operators, which allow us to connect its action with the twist by the cyclotomic character on the Galois representation side, as expressed in Theorem 13.15, whose statement we report here for convenience.

Theorem (C). Let $f \in H^i(X, \omega^{(k_1,k_2)}) \otimes \overline{\mathbb{Q}}_p$ be an eigenform for the Hecke algebra $\mathcal{H} := C_c^0(\mathrm{GSp}_4(\mathbb{A}_f)//K)$, being $i = 0$ or $i = 1$; let ρ_f be its attached Galois representation (see [11, Theorem 5.3.1]). Let $g := \vartheta(f) \in H^i(\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}, \mathfrak{F}^{(k_1+2,k_2)}) \otimes \overline{\mathbb{Q}}_p$, where $\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$ denotes $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$ minus the boundary. Assume $g \neq 0$. Then, g is still an eigenform for the prime-to-p Hecke algebra $\mathcal{H}^p$; moreover, there exists a Galois representation ρ_g attached to its Hecke eigensystem, which coincides with the cyclotomic twist of f (in other words, $\rho_g \cong \rho_f \otimes \omega$, being ω the cyclotomic character).

4. Outline

The material is organized as follows. Chapter 11 deals with the necessary preliminaries. Its heart consists of Section 11.3, containing the results about enriched p-divisible groups and their deformations that will be needed to construct the infinitesimal actions and the corresponding theta operators in the subsequent chapters. Then, Chapter 12 and Chapter 13 construct the ϑ operator in the two settings that we address, namely the μ-ordinary locus of PEL-type Shimura varieties, and the p-rank ≥ 1 locus of the Siegel threefold.

CHAPTER 11

Infinitesimal actions on enriched p-divisible groups

1. Torsors and representations

This section is devoted to introduce some notations and to recall some basic facts about actions and torsors.

Let G be a group scheme over a base scheme S, and let $\pi : \mathcal{T} \to S$ be a left G-torsor over S. For every left representation V of G over S (i.e., V is a finite locally free $\mathcal{O}_S$-module, endowed with an action $\rho : G \to \mathrm{GL}(V)$), one can form the S-scheme $\mathcal{T}^\rho := \mathcal{T} \times^G V = \{(x, v) : x \in \mathcal{T}, v \in V\}/\{(x, v) \sim (gx, gv) : g \in G\}$; in other words, $\mathcal{T}^\rho$ is the quotient of the S-scheme $\mathcal{T} \times_S V$ under the diagonal action of G.

Proposition 11.1. The S-scheme $\mathcal{T}^\rho$ is a vector bundle; as a finite, locally-free $\mathcal{O}_S$-module, it corresponds to the sheaf of G-equivariant morphisms of S-schemes $\mathcal{T} \to V$.

Proof. Let $\psi : G \xrightarrow{\sim} \mathcal{T}$ be a trivialization of the torsor $\mathcal{T}$ over some cover $S' \to S$. Then, $f_\psi : V \to \mathcal{T}^\rho, v \mapsto (\psi(1_G), v)$ and $g_\psi : \mathcal{T}^\rho \to V$, $(x, v) \mapsto (\psi^{-1}(x)v)$ are clearly inverse to each other, and hence $\mathcal{T}^\rho$ gets identified with V over S'. In particular, $\mathcal{T}^\rho$ is a vector bundle. Given a G-equivariant morphism $h : \mathcal{T} \to V$ over S', then the element $(\psi(1_G), h(\psi(1_G))) \in \mathcal{T}^\rho(S')$ is independent of the trivialization ψ chosen, and hence h canonically corresponds to a section of $\mathcal{T}^\rho$ over S'. It is easy to see that, conversely, every section of $\mathcal{T}^\rho$ determines a unique equivariant morphism $\mathcal{T} \to V$. □

The proof of the proposition emphasises that $\mathcal{T}^\rho$ can be thought of as a $\mathcal{T}$-twisted version of V, in the sense that any given trivialisation of $\mathcal{T}$ induces an isomorphism $V \cong \mathcal{T}^\rho$. This also emerges very clearly from the two examples given below.

Example 11.2. Let $\rho : G \to \mathbb{G}_{m,S}$ be a character of the group G. In this case, $\mathcal{T}^\rho$ is a line bundle, and it coincides with $(\pi_*\mathcal{O}_\mathcal{T})[-\rho]$, i.e. the invertible $\mathcal{O}_S$-subsheaf of $\pi_*\mathcal{O}_\mathcal{T}$ consisting of those functions on $\mathcal{T}$ on which G acts via its character $-\rho$.

Example 11.3. Suppose $G = \mathrm{GL}_2$, let W be a rank-2 vector bundle on S, and let $\mathcal{T}$ be the corresponding left G-torsor $\mathrm{Isom}(W, \mathcal{O}_S^2)$. Suppose we are given a dominant weight $(k_1, k_2) \in \mathbb{Z}^2$ for G (the dominance condition is $k_1 \geq k_2$), and let V_{k_1,k_2} be the irreducible representation of highest weight (k_1, k_2) of GL_2, i.e. $V_{k_1,k_2} = \mathrm{Sym}^{k_1-k_2}(\mathcal{O}_S^2) \otimes \det^{k_2}(\mathcal{O}_S^2)$ with the obvious canonical left action of $\mathrm{GL}_2 = \mathrm{Aut}(\mathcal{O}_S^2)$. Then, we have that $\mathcal{T}^\rho \cong \mathrm{Sym}^{k_1-k_2}(W) \otimes \det^{k_2}(W)$.

2. PEL data

In this section, we establish all necessary notation and assumptions that we will need to deal with PEL-type Shimura varieties, following Chapter 1 and 2 of [8], to which we refer for any detail or proof.

We recall that a *PEL datum* is given by:

(a) a finite-dimensional semisimple $\mathbb{Q}$-algebra B carrying a positive involution $* : B \to B$;
(b) a finitely generated B-module V, endowed with an alternating bilinear hermitian $\mathbb{Q}(1)$-valued pairing $\langle -, - \rangle$;
(c) an $\mathbb{R}$-algebra homomorphism $h_0 : \mathbb{C} \to \mathrm{End}_{\mathcal{O}\otimes_{\mathbb{Z}}\mathbb{R}}(L \otimes_{\mathbb{Z}} \mathbb{R})$ that satisfies the following conditions:
  (i) $\langle h_0(z)x, y \rangle = \langle x, h_0(\overline{z})y \rangle$
  (ii) the symmetric $\mathbb{R}$-bilinear pairing $(2\pi i)^{-1}\langle -, h_0(i) \cdot - \rangle$ on $V \otimes_{\mathbb{Q}} \mathbb{R}$ is positive definite.

We assume that the semisimple algebra B has no factor of type D (in other words, we will only focus on PEL-type Shimura varieties of type A and C).The morphism h_0 from point (c) can be seen as a morphism $\mathbb{S} \to G_{\mathbb{R}}$, where $\mathbb{R} := \mathrm{Res}_{\mathbb{C}/\mathbb{R}}(\mathbb{G}_{m,\mathbb{C}})$, and it defines a complex structure on $V_{\mathbb{R}}$. As a consequence, $V_{\mathbb{C}}$ decomposes as a direct sum $V_{\mathbb{C}} = V_0 \oplus \overline{V_0}$, where the element $i \in \mathbb{S}(\mathbb{R}) = \mathbb{C}^{\times}$ acts as multiplication by i (resp. by $-i$) on V_0 (resp. $\overline{V_0}$). The isomorphism class of V_0 as a $B \otimes_{\mathbb{Q}} \mathbb{C}$-module is a finite extention of $\mathbb{Q}$ that we name the *reflex field* of the PEL datum, and will be denoted by E.

Fix now a prime p. An *integral PEL datum with good reduction at* p is a PEL datum (B, V, h_0) such that B is unramified at p, together with:

(d) an order $\mathcal{O}_B$ inside B, invariant under $*$, and maximal at p;
(e) lattice $L \leq V$ stable under $\mathcal{O}_B$, such that the pairing $\langle -, - \rangle$ on V induces a perfect pairing $L_p \otimes L_p \to \mathbb{Z}_p(1)$.

To such an integral PEL datum, we can associate a smooth connected reductive group $G \to \mathrm{Spec}(\mathbb{Z}_{(p)})$, which is defined as the algebraic group of similitudes of the lattice L; in other words, for every $\mathbb{Z}_{(p)}$-algebra R, we have

$$G(R) = \{(g, \alpha) \in \mathrm{GL}(L \otimes R) \times \mathrm{GL}(\mathbb{Q}(1)) : \langle gx, gy \rangle = \alpha \langle x, y \rangle\}$$

Let us now fix the integral PEL datum $\mathcal{D} = (B, V, h_0, \mathcal{O}_B, L)$ with good reduction at p, and choose a neat compact open subgroup $K^p \leq G(\mathbb{A}_f^p)$. We will set $K_p = G(K_p)$, and $K := K_p K^p \leq G(\mathbb{A})$ (in other words, we choose hyperspecial level at p). We can define a moduli problem $F_{\mathcal{D},K}$ by attaching, to each $\mathcal{O}_{E,(p)}$-algebra R, the isomorphism classes of $\mathcal{D}$-enriched abelian schemes over R with level K structure, which is to say abelian schemes A/R together with the following additional structures:

(1) a prime-to-p polarization $\lambda : A \to A^{\vee}$;
(2) an action $\iota : \mathcal{O} \to \mathrm{End}_R(A)$, satisfying Rosati condition (which is a compatibility condition between ι and λ), and Kottwitz's determinant condition (which requires that the $\mathcal{O}_E$-action induced by ι on the finite locally free R-module $\mathrm{Lie}_{A/R}$ has the same determinant, in a suitable sense, as the B-action on V_0);
(3) a K^p-level structure, compatible with λ and ι in a suitable sense.

THEOREM 11.4. The moduli problem above is representable by a smooth, quasi-projective scheme $Y_{\mathcal{D},K} \to \mathrm{Spec}(\mathcal{O}_{E,(p)})$, which we will name the *level-K integral PEL-type Shimura variety attached to the datum $\mathcal{D}$.*

3. Enriched p-divisible groups and their extensions

Let us retain the notation and assumptions introduced in Section 11.2. The Shimura variety $Y_{\mathcal{D},K}$ carries a universal $\mathcal{D}$-enriched abelian scheme $A \to \operatorname{Spec}(\mathcal{O}_{E,(p)})$ of level K. Its p-divisible group $G := A[p^\infty]$ is an $\mathcal{O}_B$-enriched polarized p-divisible group, in the following sense.

Definition 11.5. Given a ring R, an $\mathcal{O}_B$-enriched polarized p-divisible group on R (or an $\mathcal{O}_B$-enriched Barsotti-Tate group on R) is a p-divisible group G over R together with:

(1) an $\mathcal{O}_B$-action, i.e. a morphism $\iota : \mathcal{O}_B \to \operatorname{End}_R(G)$;
(2) a polarization λ, which is to say an $\mathcal{O}_B$-linear isomorphism $\lambda : G \xrightarrow{\sim} G^\vee$ such that $\lambda^\vee = -\lambda$; which can equivalently be described as a perfect alternating pairing $T_pG \times T_pG \to T_p\mu_{p^\infty}$.

The study of such enriched p-divisible groups and their extensions is the subject of the present section.

Given a $\mathbb{Z}_p$-lattice $\mathfrak{L}$ carrying an $\mathcal{O}_B$-action, the p-divisible group $(\mathfrak{L}^\vee \otimes \mu_{p^\infty}) \oplus (\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p)$, considered with its standard polarization, is a $\mathcal{O}_B$-enriched polarized p-divisible group: we denote it $\mathbb{X}_L^{\mathrm{ord}}$ (where the ord superscript stands for "ordinary").

Lemma 11.6. The following hold:

(a) $\mathfrak{L} \cong \mathfrak{L}'$ if and only if $\mathbb{X}_{\mathfrak{L}}^{\mathrm{ord}} \cong \mathbb{X}_{\mathfrak{L}'}^{\mathrm{ord}}$
(b) if k is a separably closed field of characteristic p, all ordinary BT-groups with $\mathcal{D}$-structure over k are isomorphic to $\mathbb{X}_{\mathfrak{L}}^{\mathrm{ord}}$ for some $\mathfrak{L}$.

Proof. We remark that $\mathfrak{L}$ can be recovered as the Tate module of $(\mathbb{X}_{\mathfrak{L}}^{\mathrm{ord}})^{\text{ét}}$, from which it follows that any isomorphism $\mathbb{X}_{\mathfrak{L}}^{\mathrm{ord}} \cong \mathbb{X}_{\mathfrak{L}'}^{\mathrm{ord}}$ induces an isomorphism $\mathfrak{L} \cong \mathfrak{L}'$, which proves (a). Similarly, if k is a separably closed field, an ordinary (non-enriched) p-divisible group G over k is isomorphic, by definition of being ordinary, to $(T_p((G^{\mathrm{mult}})^\vee)^\vee \otimes \mu_{p^\infty}) \oplus T_p(G^{\text{ét}} \otimes \mathbb{Q}_p/\mathbb{Z}_p)$. If G is an $\mathcal{O}_B$-enriched polarized p-divisible group, the polarization on G identifies $T_p((G^{\mathrm{mult}})^\vee)$ with the dual of $T_p(G^{\text{ét}})$, the $\mathcal{O}_B$-action on G induces an $\mathcal{O}_B$-action on $\mathfrak{L} := T_p(G^{\text{ét}}) \cong (T_p((G^{\mathrm{mult}})^\vee))^\vee$, and G becomes isomorphic, as an $\mathcal{O}_B$-enriched polarized p-divisible group, to $\mathbb{X}_{\mathfrak{L}}^{\mathrm{ord}}$. $\square$

If $\operatorname{Nilp}_{\mathcal{O}_E}$ is the site of $\mathcal{O}_E$-algebras on which p is nilpotent, we denote $\mathrm{BT}_{\mathcal{O}_B,*} \to \operatorname{Nilp}_{\mathcal{O}_E}$ the stack attaching to each $R \in \operatorname{Nilp}_{\mathcal{O}_E}$ the category of $\mathcal{O}_B$-enriched polarized Barsotti-Tate groups on R. Above $\mathrm{BT}_{\mathcal{O}_B,*}$, we can form the the stack $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*} \to \mathrm{BT}_{\mathcal{O}_B,*}$ of $\mathcal{O}_B$-enriched polarized p-divisible groups G, together with an injection $i : \mathfrak{L}^\vee \otimes \mu_{p^\infty} \hookrightarrow G$ and a projection $\pi : G \twoheadrightarrow \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ whose kernel is p-divisible, such that i and π are dual to each other with respect to the polarization of G. A p-divisible group $G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ consequently comes with a filtration $0 = \operatorname{Fil}_0 G \subseteq \operatorname{Fil}_1 G \subseteq \operatorname{Fil}_2 G \subseteq \operatorname{Fil}_3 G = G$. If we denote $\operatorname{gr}_i(G) := \operatorname{Fil}_{i+1}(G)/\operatorname{Fil}_i(G)$, we have that the two extreme graded pieces $\operatorname{gr}_0(G)$ and $\operatorname{gr}_2(G)$ are trivialized, and are respectively multiplicative and étale, while the middle graded piece $\operatorname{gr}_1(G)$ is not. When $\ker(\pi) = \operatorname{Im}(i)$, or equivalently when $\operatorname{gr}_1(G) = 0$, G is an extension of $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ by $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$: we will denote $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*} \subseteq \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ the sub-stack consisting of extensions of $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ by $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$. For each $R \in \operatorname{Nilp}_{\mathcal{O}_E}$, we have that $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ is an abelian group, the operation being given by Baer sum of extensions.

Remark 11.7. A p-divisible group $G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ can be seen as an extension in two ways:

$$0 \to \mathrm{Fil}_1(G) \to G \to G/\mathrm{Fil}_1(G) \to 0 \qquad \text{where } \mathrm{Fil}_1(G) = \mathfrak{L}^\vee \otimes \mu_{p^\infty}$$
$$0 \to \mathrm{Fil}_2(G) \to G \to G/\mathrm{Fil}_2(G) \to 0 \qquad \text{where } G/\mathrm{Fil}_2(G) = \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$$

The two stacks $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ and $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ we have just introduced have the additional structures summarized by the following propositions.

Proposition 11.8. For each $R \in \mathrm{Nilp}_{\mathcal{O}_E}$, $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ carries a $\mathbb{Z}_p^\times$-action. On the subset $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R) \subset \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, this action extends to a $\mathbb{Z}_p$-module structure. Moreover, $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ and $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ both carry a left action by $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$.

Proof. Given an endomorphism $f \in \mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$, and a p-divisible group $G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, let us first denote by $(f,0)\cdot E$ and $(0,f)\cdot E$ the two p-divisible groups with $\mathcal{O}_B$-structure obtained by pushing out E along the morphism $f \in \mathrm{End}(\mathfrak{L}^\vee \otimes \mu_{p^\infty})$, and pulling it back along $f^\vee \in \mathrm{End}(\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p)$, respectively. The two operations are dual to each other, meaning that $((f,0)\cdot E)^\vee \cong (0,f)\cdot E^\vee$, and $((0,f)\cdot E)^\vee \cong (f,0)\cdot E^\vee$. Given two endomorphisms $f,g \in \mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$, we synthetically write (f,g) for $(f,0)\cdot(0,g)\cdot E \cong (0,g)\cdot(f,0)\cdot E$.

Let us now first construct the $\mathbb{Z}_p$-action on $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$: given $a \in \mathbb{Z}_p$ and $E \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, we have to define aE. Let us first observe that $(a,0)E$ and $(0,a)E$ fit into the following commutative diagram.

$$\begin{tikzcd}
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r, "i'"] \arrow[d, equal] & (0,a)\cdot E \arrow[r, "\pi'"] \arrow[d, "f"] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] \arrow[d, "a"] & 0 \\
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r, "i"] \arrow[d, "a"] & E \arrow[r, "\pi"] \arrow[d, "g"] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] \arrow[d, equal] & 0 \\
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r, "i''"] & (a,0)\cdot E \arrow[r, "\pi''"] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] & 0
\end{tikzcd}$$

It follows from the universal properties of pushout and pullback that there exists a canonical morphism ψ from $(a,0)E$ to $(0,a)E$ such that $\pi'\psi g = \pi$, $\pi'\psi i'' = 0$, $f\psi g = a$, $f\psi i'' = i$. We have that ψ restricts to the identity on both $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$ and on $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$, and consequently establishes an isomorphism between the two extensions. We can thus set $aE := (a,0)E \cong (0,a)E$; from the polarization $\lambda : E \xrightarrow{\sim} E^\vee$, one obtains a polarization $\lambda' : aE \xrightarrow{\sim} a(E^\vee) = (a,0)E^\vee = ((0,a)E)^\vee = (aE)^\vee$.

The $\mathbb{Z}_p^\times$-action on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ simply consists of multiplying by the invertible scalar a the polarization carried by the p-divisible group. We omit the verification that this operation is compatible with the $\mathbb{Z}_p$-action on $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ described above.

We are now only left to show the existence of the $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$-action on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$. Given an endomorphisms $f \in \mathrm{End}(\mathfrak{L}^\vee \otimes \mu_{p^\infty})$, we define this action as $f \cdot E := (f,f)\cdot E$. If we take the polarization $\lambda : E \xrightarrow{\sim} E^\vee$ and apply $f \cdot -$ on both sides, we get $\lambda' : f\cdot E \xrightarrow{\sim} f \cdot E^\vee = (0,f)\cdot(f,0)\cdot E^\vee \cong ((f,0)\cdot(0,f)\cdot E)^\vee = (f\cdot E)^\vee$. As a result, $f \cdot E$ also carries a canonical polarization induced from that of E, and hence $f \cdot E \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$. $\square$

Remark 11.9. Regarding the actions introduced in Proposition 11.8 on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ and $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, we make the following remarks:

(a) Given $a \in \mathbb{Z}_p$, it is clear that a can also be thought of as an endomorphism of $\mathfrak{L}^\vee$. However, given $E \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, the extensions aE and $a \cdot E$ that we respectively obtain by letting $a \in \mathbb{Z}_p$ and $a \in \mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$ act on E do not coincide; instead, the proof of the previous proposition shows that $a \cdot E$ coincides with $a^2 E$.

(b) Since $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ and $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ are modules for $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$, they are acted upon by $\mathrm{Aut}(\mathfrak{L})$, where, given $f \in \mathrm{Aut}(\mathfrak{L})$ and $G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, we define $f \cdot E := g \cdot E$, where $g := f^{\vee,-1} \in \mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$.

We will now exhibit a way of constructing elements of $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$. Let us fix $R \in \mathrm{Nilp}_{\mathcal{O}_E}$, and consider a $\mathbb{Z}_p$-linear morphism $Z : \mathfrak{L} \otimes_{\mathcal{O}_B} \mathfrak{L} \to \mu_{p^\infty}$ over R. This can be seen as a Hermitian bilinear μ_{p^∞}-valued pairing $\langle -, - \rangle_Z$ on $\mathfrak{L}$, and it induces an $\mathcal{O}_B$-linear morphism $\varphi_Z : \mathfrak{L} \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$, $x \mapsto \langle x, - \rangle_Z$. We can construct an extension E_Z by taking the short exact sequence $0 \to \mathbb{Z}_p \to \mathbb{Q}_p \to \mathbb{Q}_p/\mathbb{Z}_p \to 0$, tensoring with $\mathfrak{L}$, and taking the pushout along φ_Z:

$$
\begin{tikzcd}
0 \arrow[r] & \mathfrak{L} \arrow[r] \arrow[d, "\varphi_Z"] & \mathfrak{L} \otimes_{\mathbb{Z}_p} \mathbb{Q}_p \arrow[r] \arrow[d] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] \arrow[d, equal] & 0 \\
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r] & E_Z \arrow[r] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] & 0
\end{tikzcd}
$$

Proposition 11.10. Assume that Z is a symmetric form, then the extension E_Z carries a canonical polarization, and can consequently be seen as an $\mathcal{O}_B$-enriched polarized p-divisible group over R. The assignment $Z \mapsto E_Z$ thus defines a $\mathbb{Z}_p$-linear ($\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$)-equivariant map $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty} \to \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, where $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee)$ denotes the $\mathbb{Z}_p$-module of bilinear Hermitian symmetric forms on $\mathfrak{L}$.

Proof. To prove the result, we will exhibit a collection perfect alternating pairings $(\cdot,\cdot) : E_Z[p^n] \times E_Z[p^n] \to \mu_{p^n}$ that are compatible as n varies, and thus provide E_Z with a polarization. It follows from the definition that E_Z admits the following description:

$$E_Z = \frac{\mathrm{Hom}(\mathfrak{L}, \mu_{p^\infty}) \oplus (\mathfrak{L} \otimes \mathbb{Q}_p)}{\{(\langle x, \cdot \rangle_Z, -x \otimes 1) : x \in \mathfrak{L}\}}.$$

Given two elements $w_i := (f_i, x_i \otimes m_i) \in E_Z$, for $i = 1, 2$, suppose that they lie in the p^n-torsion subgroups, which is equivalent to requiring that f_i, x_i and m_i can be chosen so that:

(i) $m_i \in p^{-n}\mathbb{Z}_p \subset \mathbb{Q}_p$;

(ii) $f_i : \mathfrak{L} \to \mu_{p^\infty}$ becomes equal to $\langle x_i, \cdot \rangle_Z^{-m_i}$ after raising to the p^n-power.

We define the alternating pairing to be

$$(w_1, w_2)_Z := f_1(x_2)^{p^n m_2} f_2(x_1)^{-p^n m_1}.$$

To check that this is well-defined, one must verify that if either of the w_1 or w_2 has the form $(\langle x, \cdot \rangle_Z, -x \otimes 1)$, we get $(w_1, w_2) = 1$. If we assume, without losing generality, that w_1 has this form, we get:

$$(w_1, w_2)_Z = [\langle x, x_2 \rangle_Z]^{p^n m_2} [f_2(-x)]^{-p^n} = [\langle x, x_2 \rangle_Z]^{p^n m_2} \left[\langle x_2, -x \rangle_Z^{-m_2}\right]^{-p^n} = 1.$$

We have thus exhibited the polarization on E_Z. We leave to the reader the verification of the fact that the assignment $Z \in \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty} \mapsto E_Z \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ has the linearity properties listed in the statement. □

We can characterize the image of the map $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty} \to \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, $Z \mapsto E_Z$ as follows.

Lemma 11.11. For an extension $E \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$, denoting ι and π the inclusion $\mathfrak{L}^\vee \otimes \mu_{p^\infty} \to E$ and the projection $E \to \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ respectively, the following are equivalent.

(a) E is in the image of the map $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty} \to \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, $Z \mapsto E_Z$.

(b) $E \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)^{\mathrm{tors}}$, meaning that $p^n E$ is a trivial extension for n sufficiently large.

(c) There exists $j : \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \to E$ such that $\pi j = p^n$ for some n.

(d) There exists $q : E \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$ such that $qi = p^n$ for some n.

Proof. Since $\mathfrak{L}^\vee \otimes \mu_{p^\infty}(R)$ is a p-torsion group, and $Z \mapsto E_Z$ is $\mathbb{Z}_p$-linear, it is clear that (a) implies (b). Let us now prove that (b) implies (c). The extension $p^n E$ is the pushout of E along the morphism $p^n : \mathfrak{L}^\vee \otimes \mu_{p^\infty} \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$:

$$\begin{tikzcd}
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r, "\iota"] \arrow[d, "p^n"] & E \arrow[r, "\pi"] \arrow[d] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] \arrow[d, equal] & 0 \\
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r, "\iota'"] & p^n E \arrow[r, "\pi'"] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] & 0
\end{tikzcd}$$

The existence of a splitting $q' : p^n E \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$ of ι' yields the existence of a morphism $q : E \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$ such that $q\iota = p^n$, which proves (c). That (c) and (d) are equivalent is clear. Finally, let us show that (d) implies (a). Let $j : \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \to E$ be a morphism such that $\pi j = p^n$. Then, $p^{-n} j : \mathfrak{L} \otimes \mathbb{Q}_p \to E$ induces the following morphisms of extensions

$$\begin{tikzcd}
0 \arrow[r] & \mathfrak{L} \arrow[r] \arrow[d] & \mathfrak{L} \otimes \mathbb{Q}_p \arrow[r] \arrow[d, "p^{-n}j"] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] \arrow[d, equal] & 0 \\
0 \arrow[r] & \mathfrak{L}^\vee \otimes \mu_{p^\infty} \arrow[r] & E \arrow[r] & \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p \arrow[r] & 0
\end{tikzcd}$$

which shows that $E = E_Z$ for some $Z : \mathfrak{L} \to \mathfrak{L}^\vee \otimes \mu_{p^\infty}$. □

We are now ready to construct an action of $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, which will be the ultimate source of the ϑ operators constructed in this work.

Proposition 11.12. There is a natural action

$$\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*} \times \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*} \to \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}, \quad (E, G) \mapsto E + G$$

which is equivariant with respect to the actions of $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$ and $\mathbb{Z}_p^\times$ that we have on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ and on $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$. Moreover, for each G and E, we have the following properties:

(a) there exist canonical isomorphisms $\mathrm{Fil}_{\bullet+d}(E+G)/\mathrm{Fil}_\bullet(E+G) \cong \mathrm{Fil}_{\bullet+d}(G)/\mathrm{Fil}_\bullet(G)$ for $d = 1, 2$;

(b) the operation has good functoriality properties: if two morphisms $f : G \to G'$ and $g : E \to E'$ are given respecting the filtrations, if they are equal once restricted to $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$ and to $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$, then they induce a morphism $G+E \to G'+E'$, which coincides with f when restricted to $\mathrm{Fil}_{\bullet+d}(E+G)/\mathrm{Fil}_\bullet(E+G) \cong \mathrm{Fil}_{\bullet+d}(G)/\mathrm{Fil}_\bullet(G)$ for $d = 1, 2$;

(c) the operation is compatible with Baer sums in $\mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, meaning that $E_1+(E_2+G) = (E_1+E_2)+G$;

(d) if E is the trivial extension, then $E+G \cong G$;

(e) $(E+G)^\vee \cong E^\vee + G^\vee$.

Proof. Recall that a p-divisible group $G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ carries a 3-step filtration, and can consequently be seen as an extension in two different ways, as Remark 11.7 explains. Correspondingly, there are two possible ways of defining $E+G$:

(1) One pulls back E along the projection of $G/\mathrm{Fil}_1(G)$ to $\mathrm{gr}_2(G) = \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$: in this way, one gets an extension E' of $G/\mathrm{Fil}_1(G)$ by $\mathrm{Fil}_1(G) = \mathfrak{L}^\vee \otimes \mu_{p^\infty}$, and $E+G$ can be now defined as the Baer sum of E' and G inside $\mathrm{Ext}^1(G/\mathrm{Fil}_1(G), \mathrm{Fil}_1(G))$.

(2) One pushes out E along the inclusion of $\mathrm{Fil}_1(G) = \mathfrak{L}^\vee \otimes \mu_{p^\infty}$ inside $\mathrm{Fil}_2 G$: in this way, one gets an extension E'' of $G/\mathrm{Fil}_2(G) = \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ by $\mathrm{Fil}_2(G)$, and $E+G$ can be now defined as the Baer sum of E'' and G inside $\mathrm{Ext}^1(G/\mathrm{Fil}_2(G), \mathrm{Fil}_2(G))$.

The two definitions are readily proved to give the same result (the verification is carried out, for example, in [5, §9.3]), which is a filtered p-divisible group with $\mathcal{O}_B$-structure. Construction (1) makes it evident that the operation $G \mapsto E+G$ does not alter $\mathrm{Fil}_1(G)$ and $G/\mathrm{Fil}_1(G)$; the same is true for $\mathrm{Fil}_2(G)$ and $G/\mathrm{Fil}_2(G)$ if one looks at construction (2). This can be summarized by saying that there exist canonical identifications $\mathrm{Fil}_{\bullet+d}(E+G)/\mathrm{Fil}_\bullet(E+G) \cong \mathrm{Fil}_{\bullet+d}(G)/\mathrm{Fil}_\bullet(G)$ for $d = 1, 2$, which proves that property (a) is satisfied. We omit the verification of the remaining properties of the operation, for which we refer again to [5], as well as the verification of the equivariance with respect to the action by $\mathbb{Z}_p^\times$ and $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$.

We remark that the resulting filtered p-divisible group $E+G$ is actually en element of $\mathrm{BT}_{\mathcal{O}_B,*}(\mathfrak{L})$: the $\mathcal{O}_B$-action is inherited from the one we have on E and G; the trivializations of its gr_0 and gr_2 pieces are those inherited from G; finally, the polarization on $E+G$ is the one canonically induced by those on E and G: more precisely, if we combine together $\lambda_E : E \xrightarrow{\sim} E^\vee$ with $\lambda_G : G \xrightarrow{\sim} G^\vee$, one gets, thanks to (b) an isomorphism $\lambda_{E+G} : E+G \xrightarrow{\sim} E^\vee + G^\vee \cong (E+G)^\vee$ (where the last identification is property (e)). □

4. Actions by infinitesimal multiplicative groups

4.1. Iwasawa algebras. In this subsection, we recall some introductory notions about measures and distributions on profinite groups; our reference is [9].

Let R be a p-adically complete $\mathbb{Z}_p$-algebra. Given a profinite group G, one can define its complete group algebra $R[[G]] := \varprojlim_H R[G/H]$, where H ranges among all open normal subgroups of H, and G/H consequently ranges among all finite discrete quotients of H. This is also known as the Iwasawa algebra of G.

Proposition 11.13. The elements of $R[[G]]$ correspond to the R-valued *measures* on G, where a R-valued measure on G is a rule μ attaching to each closed and open subset E of G a scalar $\mu(E) \in R$, in such a way that $\mu(\emptyset) = 0$ and $\mu(E_1 \sqcup \ldots \sqcup E_n) = \mu(E_1) + \ldots + \mu(E_n)$.

Proof. Let $x \in R[[G]]$, and let E be a closed and open subset of G. Then, $E = Hg_1 \sqcup \ldots \sqcup Hg_k$ is a finite union of cosets of a certain open normal subgroup H of G. The image of $x \in R[[G]]$ in $R[G/H]$ has the form $\sum_{C \in G/H} a_C C$ for some coefficients $a_C \in R$. Then, defining $\mu(E) := \sum_{i=1}^{k} a_{Hg_i}$ gives rise to the measure μ attached to x.

Conversely, suppose μ is a measure on G. Then, for each compact normal subgroup H, one can form the element $\sum_{C \in G/H} \mu(C)C \in R[G/H]$. Thanks to the additivity of μ, these elements are compatible as H varies, and can be packed together to form an element of $R[[G]]$. $\square$

For each element $z \in G$, one can consider the atomic measure μ_z assigning 1 or 0 to E depending on whether $z \in E$ or $z \notin E$. Atomic measures are clearly dense in $R[[G]]$.

The product operation in $R[[G]]$ is given by convolution, which is the unique R-bilinear product on $R[[G]]$ such that $\mu_x \mu_y = \mu_{xy}$; this is only commutative when G is. The unit element 1 for this product is μ_e, being e the identity element of the group G, and the structure morphism $R \to R[[G]]$ sends $a \in R$ to the measure $a\mu_e$. In fact, $R[[G]]$ is not only an algebra, but a co-commutative formal Hopf algebra: the counit is given by the morphism $R[[G]] \to R$ that sends $\mu_x \mapsto 1$ for all x, while the coproduct is the unique algebra homomorphism $\Delta : R[[G]] \to R[[G]] \hat{\otimes} R[[G]]$ taking $\mu_x \mapsto \mu_x \otimes \mu_x$.

Remark 11.14. The atomic measures μ_x, varying $x \in G$, can be recovered as the *group-like* elements of $R[[G]]$, i.e. those elements $\mu \in R[[G]]$ such that $\Delta(\mu) = \mu \otimes \mu$.

Given G a profinite group, there is a second important formal Hopf algebra over R one can form, that is $\mathrm{Cont}(G, R)$. The characteristic functions χ_{aH} of the cosets of the open normal subgroups H of G are dense in $\mathrm{Cont}(G, R)$. Sums and products are computed pointwise; in particular, $\mathrm{Cont}(G, R)$ is a commutative formal Hopf algebra. Given $f \in \mathrm{Cont}(G, R)$, the comultiplication $\Delta(f)$ is computed as $\Delta(f)(g_1, g_2) = f(g_1 g_2)$; the counit is simply given by evaluating at $e \in G$.

Remark 11.15. The *additive* functions $G \to R$ (i.e., those functions f satisfying $f(g_1 g_2) = f(g_1) + f(g_2)$) can be recovered as the primitive elements of the formal Hopf algebra $\mathrm{Cont}(G, R)$, i.e. those elements f such that $\Delta(f) = 1 \otimes f + f \otimes 1$, while the *characters* $G \to R$ are the group-like elements of $\mathrm{Cont}(G, R)$, i.e. those elements f such that $\Delta(f) = f \otimes f$.

Given $\mu \in R[[G]]$ and $f \in \mathrm{Cont}(G, R)$, one can compute the integral $\int f d\mu$. When f is locally constant, i.e. when it factors through some discrete quotient G/H, this is just the sum $\sum_{E \in G/H} f(E)\mu(E)$; by density, one then easily extends the definition to the whole $\mathrm{Cont}(G, \mathbb{Z}_p)$.

Proposition 11.16. The pairing $\int$: $\mathrm{Cont}(G, R) \times R[[G]] \to R$ described above is a continuous perfect bilinear pairing of topological formal Hopf algebras, meaning that it induces topological isomorphisms of formal Hopf algebras $\mathrm{Cont}(G, R) \cong \mathrm{Hom}_R(R[[G]], R)$ and $R[[G]] \cong \mathrm{Hom}_R(\mathrm{Cont}(G, R), R)$, where Hom_R stands for continuous R-linear homomorphisms.

Proof. Let us fix $n \geq 1$. We remark that $\mathrm{Cont}(G, R/p^n) \cong \mathrm{Cont}(\varprojlim_H G/H, R/p^n) \cong \varinjlim_H \mathrm{Cont}(G/H, R/p^n)$; the last isomorphism follows from the fact that R/p^n is discrete, and hence every continuous function from G to R/p^n must factor through some discrete quotient of G. Following a similar argument, one has that $\mathrm{Hom}_R(R[[G]], R/p^n) \cong \varinjlim_H \mathrm{Hom}_R(R[G/H], R/p^n)$. But now, $\mathrm{Hom}_R(R[G/H], R/p^n) \cong \mathrm{Cont}(G/H, R/p^n)$ by the universal property of group algebras. We conclude that $\mathrm{Cont}(G, R/p^n) \cong \mathrm{Hom}_R(R[[G]], R/p^n)$ for all n. Taking the inverse limit on n, one gets $\mathrm{Cont}(G, R) \cong \mathrm{Hom}_R(R[[G]], R)$.

The other isomorphism, namely $R[[G]] \cong \mathrm{Hom}_R(\mathrm{Cont}(G, R), R)$ by proving, by similar arguments, that both of them coincide with $\varprojlim_H \varprojlim_n \mathrm{Cont}(G/H, R/p^n)^\vee$. $\square$

4.2. Equivalent perspectives on μ_{p^∞}-actions. In this subsection, we will present equivalent ways of describing an action by the infinitesimal group μ_{p^∞} on a formal scheme.

Proposition 11.17. Given a $\mathbb{Z}_p$-lattice M, and $R \in \mathrm{Nilp}_{\mathbb{Z}_p}$, then $R[[M^\vee]]$ is canonically isomorphic to the formal Hopf algebra of the formal group $M \otimes \mu_{p^\infty}$ over $\mathrm{Spf}(R)$.

Proof. Given $A \in \mathrm{Nilp}_R$, an A-point of $\mathrm{Spf}(R[[M^\vee]])$ is a continuous algebra homomorphism $R[[M^\vee]] \to A$, which necessarily factors through some $R[M^\vee/p^n M^\vee]$ giving rise to a morphism $R[M^\vee/p^n M^\vee] \to A$, i.e. a group homomorphism $f : M^\vee/p^n M^\vee \to A^\times$. Since the domain of f is a p^n-torsion group, the image of f can only consist of p^n-roots of unity, so that f can equivalently be interpreted as a morphism $M^\vee \to \mu_{p^\infty}[p^n](A)$, i.e. an A-point of $(M \otimes \mu_{p^\infty})[p^n]$. $\square$

Proposition 11.18. Suppose $\mathfrak{X}$ is a formal group scheme over a p-adically complete base ring R. Then, the following data are equivalent:

(a) a left action by $M \otimes \mu_{p^\infty}$ on $\mathfrak{X}$;
(b) a right co-action by the formal Hopf algebra $R[[M^\vee]]$ on the sheaf of algebras $\mathcal{O}_{\mathfrak{X}}$;
(c) a left action by the formal Hopf algebra $\mathrm{Cont}(M^\vee, R)$ on the sheaf of algebras $\mathcal{O}_{\mathfrak{X}}$.

Proof. We will only prove how an action by $M \otimes \mu_{p^\infty}$ entails the algebra actions (b) and (c). Let us thus suppose that $\mathfrak{X}$ is endowed with an action of $M \otimes \mu_{p^\infty}$. We remark that, since the group $M \otimes \mu_{p^\infty}$ is infinitesimal, it acts on $\mathfrak{X}$ via universal homeomorphisms; in particular, the action preserves all open subsets of the formal scheme $\mathfrak{X}$; if f is a function defined on some open subscheme U of $\mathfrak{X}$, and $\zeta \in (M \otimes \mu_{p^\infty})(R)$, then one can define $\zeta \cdot f$ to be $f : U \to \mathbb{A}^1$ precomposed with the automorphism $U \to U$ induced by ζ^{-1}. This yields a left action by $M \otimes \mu_{p^\infty}$ on $\mathcal{O}_{\mathfrak{X}}$, which is to say a right co-module structure on $\mathcal{O}_{\mathfrak{X}}$ for the formal Hopf algebra $R[[M^\vee]]$, given by a morphism $\mathcal{O}_{\mathfrak{X}} \to \mathcal{O}_{\mathfrak{X}} \hat{\otimes} R[[M^\vee]]$. But this is the same as a left module structure on $\mathcal{O}_{\mathfrak{X}}$ by the dual formal Hopf algebra $\mathrm{Cont}(M^\vee, R)$. $\square$

So, given a section $s : U \to \mathbb{A}^1_R$ over some open affine subset $U \subseteq \mathfrak{X}$, and a continuous function $f : M^\vee \to R$, one can form $f \cdot s \in \mathcal{O}_{\mathfrak{X}}(U)$. An immediate generalization of this result is stated here below, where we replace the codomains of s and f with vector bundles over R.

Proposition 11.19. Let $\mathfrak{X} \to \mathrm{Spf}(R)$ be a formal scheme endowed with an action by the infinitesimal multiplicative group $M \otimes \mu_{p^\infty}$, as above. Let $U \subseteq R$ be a formal open subscheme. Given two vector bundles V and W over R, an element $f \in \mathrm{Cont}(M^\vee, V)$, and a function $s : U \to W$, we have that f acts on s, taking it to a function $f \cdot s : U \to V \otimes_R W$.

Proof. The action of $M \otimes \mu_{p^\infty}$ on $\mathfrak{X}$, using the point of view of Proposition 11.18(c) above, can be expressed as a morphism $p : \mathrm{Cont}(M^\vee, R) \otimes \mathcal{O}_\mathfrak{X} \to \mathcal{O}_\mathfrak{X}$. Tensoring both sides with $V \otimes W$ gives a morphism $p_{V,W} : \mathrm{Cont}(M^\vee, V) \otimes (W \otimes \mathcal{O}_\mathfrak{X}) \to (V \otimes W) \otimes \mathcal{O}_\mathfrak{X}$, which is the desired action. $\square$

For our applications, we will actually need to work in a more general setting, in which V and W, instead of just being R-vector bundles, are representations of a formal affine group scheme $G \to \mathrm{Spf}(R)$.

Proposition 11.20. Let $\mathfrak{S} \to \mathrm{Spf}(R)$ be a formal scheme, $G \to \mathrm{Spf}(R)$ an affine formal group scheme, and $\pi : \mathfrak{X} \to \mathfrak{S}$ a G-torsor. Suppose that:

(a) $\mathfrak{X} \to \mathrm{Spf}(R)$ carries an action by the infinitesimal multiplicative group $M \otimes \mu_{p^\infty}$;
(b) G acts on the lattice M;
(c) the two actions are compatible, meaning that $(g \cdot h) \cdot x$ and $g \cdot (h \cdot (g^{-1} \cdot x))$ coincide for all $g \in G$, $h \in M \otimes \mu_{p^\infty}$, $x \in \mathfrak{X}$.

Let V and W be representations of G over R. Then, given a G-equivariant continuous function $f : M^\vee \to V$, and a G-equivariant function $s : \mathfrak{X}_{|U} \to W$, we have that the function $f \cdot s : \mathfrak{X}_{|U} \to V \otimes W$ defined in Proposition 11.19 is also G-equivariant.

Proof. Let us view the $(M \otimes \mu_{p^\infty})$-action on $\mathfrak{X}$ as a module structure $p : \mathrm{Cont}(M^\vee, R) \otimes \mathcal{O}_\mathfrak{X} \to \mathcal{O}_\mathfrak{X}$ (see Proposition 11.18(c)). Assumption (c) can be expressed by saying that p is equivariant with respect to the action of G. It will thus consequently remain G-equivariant if both sides are tensored with the G-representation $V \otimes W$, which means that the morphism $p_{V,W} : \mathrm{Cont}(M^\vee, V) \otimes (W \otimes \mathcal{O}_\mathfrak{X}) \to (V \otimes W) \otimes \mathcal{O}_\mathfrak{X}$ introduced in the proof of Proposition 11.19 is also G-equivariant. But this implies, in particular, that it takes G-invariant elements of its domain to G-invariant elements of its codomain, whence the proposition follows. $\square$

4.3. Differentiating μ_{p^∞}-actions. Given a section s of $\mathcal{O}_\mathfrak{X}$ over an open affine subset $U \subseteq \mathfrak{X}$, an element $\zeta \in (M \otimes \mu_{p^\infty})(R)$, we want to construct $\zeta \cdot s$, and this can be done in two equivalent ways.

(1) we see s as a function $U \to \mathbb{A}^1$, and precompose it with the automorphism $\psi_{\zeta^{-1}} : U \to U$ given by the action of $\zeta \in (M \otimes \mu_{p^\infty})(R)$: in other words, $\zeta \cdot s := s \circ \psi_{\zeta^{-1}}$;
(2) we consider the continuous function $M^\vee \to R$, $x \mapsto \zeta^x$, which is the *group-like* element of $\mathrm{Cont}(M^\vee, R)$ corresponding to ζ, and we let it act on s adopting the point of view of Proposition 11.18(c); in other words, $\zeta \cdot s := [x \mapsto \zeta^x] \cdot s$.

The Lie algebra of $G := M \otimes \mu_{p^\infty}$ is canonically isomorphic to $M \otimes R$, the isomorphism being induced by $M \to \mathrm{Lie}_{G/R}(R)$, $m \mapsto m \otimes (1 + \varepsilon) \in (M \otimes \mu_{p^\infty})(R[\varepsilon])$, where $\varepsilon^2 = 0$. For each function $f \in \mathcal{O}_\mathfrak{X}(U)$, we have that $(m \otimes (1 + \varepsilon)) \cdot f = f + \varepsilon D_m f$, for some section $D_m f \in \mathcal{O}_\mathfrak{X}(U)$ that depends linearly on $m \in M$. The operation $f \mapsto D_m f$ can be thought of as the *derivative* of the action of μ_{p^∞} on $\mathfrak{X}$ *in the direction* m. The continuous character $M^\vee \to R[\varepsilon]$ corresponding to $\zeta = m \otimes (1 + \varepsilon)$ is $x \mapsto (1 + \varepsilon)^{\langle x, m\rangle} = 1 + \langle x, m\rangle\varepsilon$, which shows that $D_m f$ can alternatively be defined as the image of s through the action of the function $M^\vee \to R$, $x \mapsto \langle x, m\rangle$. Note that $[x \mapsto \langle x, m\rangle]$ is a primitive element of the formal Hopf algebra $\mathrm{Cont}(M^\vee, R)$, which entails that D_m satisfies the Leibniz rule $D_m(f_1 f_2) = D_m(f_1) f_2 + f_1 D_m(f_2)$.

Definition 11.21. Given $\mathfrak{X} \to \mathrm{Spf}(R)$ a formal scheme over a p-adically complete $\mathbb{Z}_p$-algebra R, endowed with a $(M \otimes \mu_{p^\infty})$-action, and given an element $m \in M$, the *differential of the action in the direction* m is the R-linear differential operator $D_m : \mathcal{O}_{\mathfrak{X}} \to \mathcal{O}_{\mathfrak{X}}$ introduced above.

Remark 11.22. The assignment $m \mapsto D_m$ is $\mathbb{Z}_p$-linear.

CHAPTER 12

Theta operators on the μ-ordinary locus

1. The μ-ordinary setting

We keep all notation from Section 11.2, where we have defined a Shimura variety $Y_{\mathcal{D},K} \to \mathrm{Spec}(\mathcal{O}_{E,(p)})$ attached to a level K (hyperspecial at p) and an integral PEL datum $\mathcal{D} = (B, V, h_0, \mathcal{O}_B, L)$, both satisfying suitable assumptions: we will denote it Y for brevity, and it carries a universal $\mathcal{D}$-enriched abelian scheme with level K, denoted $A \to Y$. Let now $X := X_{\mathcal{D},K} \to \mathrm{Spec}(\mathcal{O}_{E,(p)})$ denote a toroidal compactification of the Shimura variety; the abelian scheme $A \to Y$ extends to a semi-abelian scheme $A \to X$, and we denote by $D := X \smallsetminus Y$ the boundary divisor. The sheaf of invariant differentials of A, which we denote ω, is a vector bundle on X, i.e. a torsor for the group $M := \mathrm{GL}_n$.

1.1. The μ-ordinary locus. Let $\mathfrak{X} \to \mathrm{Spf}(\mathcal{O}_{E,p})$ be the formal completion of X at p, and let $\mathfrak{Y}$ be that of Y. Inside $\mathfrak{X}$, one can identify an open dense formal subscheme $\mathfrak{X}^{\mu-\mathrm{ord}}$, known as the *$\mu$-ordinary locus*, that can be identified as the non-vanishing locus of a number of partial Hasse invariants. Away from the boundary, the μ-ordinary locus can be characterized by the fact that the ($\mathcal{O}_B$-enriched, polarized) p-divisible group of A is (fiberwise) μ-ordinary (for the definition of a μ-ordinary enriched p-divisible group, see [10]). Over $\mathfrak{Y}^{\mu-\mathrm{ord}}$, $A[p^\infty]$ is isotrivial, and carries a slope filtration. We denote by $A[p^\infty]^\lambda$ the λ-slope piece of $A[p^\infty]$, so that, in particular, $A[p^\infty]^{\lambda=0}$ is the multiplicative part, while $A[p^\infty]^{\lambda=1}$ is the étale part. All isoclinic pieces $A[p^\infty]^\lambda$ are isotrivial, and the polarization induces a duality between $A[p^\infty]^\lambda$ and $A[p^\infty]^{1-\lambda}$.

We can thus choose isoclinic μ-ordinary p-divisible groups $\mathbb{X}^\lambda \in \mathrm{BT}_{\mathcal{O}_B,*}(\mathcal{O}_{E,(p)})$ representing the isomorphism type of $A[p^\infty]^\lambda$, together with identifications $\psi_\lambda : \mathbb{X}^\lambda \xrightarrow{\sim} \mathbb{X}^{1-\lambda}$ such that $\psi_\lambda^\vee = -\psi_{1-\lambda}$, which make $\mathbb{X} := \oplus_\lambda \mathbb{X}^\lambda$ a polarized p-divisible group with $\mathcal{O}_B$-action. At each geometric fiber of $\mathfrak{Y}^{\mu-\mathrm{ord}}$, we have that $A[p^\infty]$ is isomorphic to $\mathbb{X}$.

For $\lambda = 0$ and $\lambda = 1$, we choose $\mathbb{X}^{\lambda=0} := \mathfrak{L}^\vee \otimes \mu_{p^\infty}$ and $\mathbb{X}^{\lambda=1} := \mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ respectively, for some suitably $\mathbb{Z}_p$-lattice with $\mathcal{O}_B$-action $\mathfrak{L}$: this is possible thanks to Lemma 11.6.

The slope filtration on $A[p^\infty]$ contravariantly induces a slope filtration on the sheaf of differentials ω; we let ω^λ denote the sheaf of differentials of the graded piece $A[p^\infty]^\lambda$.

Let $\lambda_1, \ldots, \lambda_n$ denote the slopes appearing in the μ-ordinary p-divisible group $A[p^\infty]$, and let $r(\lambda_i)$ be the rank of ω^{λ_i}. Let $M^{\mathrm{gr}} \leq M$ denote the subgroup $\mathrm{GL}_{r(\lambda_1)} \times \ldots \times \mathrm{GL}_{r(\lambda_n)} \leq \mathrm{GL}_n = M$. Then, over $\mathfrak{X}^{\mu-\mathrm{ord}}$ we can construct the left M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$ of the trivializations $\oplus_\lambda \psi_\lambda : \oplus_\lambda \mathrm{gr}^\lambda(\omega) \xrightarrow{\sim} \oplus_\lambda \mathrm{gr}^\lambda(\omega_{\mathbb{X}^\lambda})$.

1.2. p-adic interpolation. Let now $\mathfrak{T}^{\mathrm{gr}}$ be the moduli space that classifies isomorphisms $\iota = \oplus_\lambda \iota_\lambda : \oplus_\lambda \mathbb{X}^\lambda \xrightarrow{\sim} \oplus_\lambda \mathrm{gr}_\lambda A[p^\infty]$: this is a $\mathfrak{M}^{\mathrm{gr}}$-torsor, where $\mathfrak{M}^{\mathrm{gr}}$ is the subgroup of $\prod_\lambda \mathrm{Aut}(\mathbb{X}_\lambda)$ consisting of those automorphisms that respect the dualities ψ_λ. We have that $\mathfrak{M}^{\mathrm{gr}}$ is a p-adic group; it will be later useful to factor it as $\mathfrak{M}^{\mathrm{gr}} = \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \times \mathfrak{M}'$, where $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$ gives the automorphism of the multiplicative and étale graded pieces, while $\mathfrak{M}'$ consists of the automorphisms of λ-graded pieces for $0 < \lambda < 1$.

If we take the differential of an automorphism in $\mathfrak{M}^{\mathrm{gr}}$, we obtain automorphisms of the sheaf of invariant differentials $\omega_{\mathbb{X}^\lambda}$, and this leads to an immersion $\mathfrak{M}^{\mathrm{gr}} \leq M^{\mathrm{gr}}$. Similarly, each trivialization of the graded pieces of $A[p^\infty]$ induces a trivialization of the graded pieces of ω, and the assignment $(\iota, \psi_1) \mapsto (d\iota, \psi_1)$ consequently defines an equivariant morphism from the $\mathfrak{M}^{\mathrm{gr}}$-torsor $\mathfrak{T}^{\mathrm{gr}}$ to the M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$, which identifies $\mathcal{T}^{\mathrm{gr}}$ as the pushforward of $\mathfrak{T}^{\mathrm{gr}}$ along the inclusion $\mathfrak{M}^{\mathrm{gr}} \leq M^{\mathrm{gr}}$. In other words, the torsor $\mathfrak{T}^{\mathrm{gr}}$ we have just constructed is a $\mathfrak{M}^{\mathrm{gr}}$-reduction of the M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$.

Given a representation $\rho : \mathfrak{M}^{\mathrm{gr}} \to \mathrm{GL}_n(R)$, for R some p-adically complete $\mathcal{O}_E$-algebra, we can thus form the vector bundle $(\mathfrak{T}^{\mathrm{gr}})^\rho$ on $\mathfrak{Y}^{\mu-\mathrm{ord}} \hat{\otimes} R$ (as explained in Section 11.1), which is the vector bundle of weight-ρ p-adic modular forms over the μ-ordinary locus, that we denote ω^ρ for brevity (see [3] for a more complete treatment of this topic).

2. The construction of ϑ

We are now ready to define the infinitesimal action on $\mathfrak{T}^{\mathrm{gr}}$ that will give rise to the ϑ operator. The acting group will turn out to be $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$.

Lemma 12.1. Let Z be a section of the formal group $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$ over some algebra $R \in \mathrm{Nilp}_{\mathcal{O}_E}$. Then, there exists a nilpotent ideal $I \leq R$ modulo which Z coincides with the identity element of the group.

Proof. Since $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$, as a p-divisible group, is (non-canonically) isomorphic to a direct sum of a certain finite number of copies of μ_{p^∞}, showing the result for μ_{p^∞} will suffice. In the case of μ_{p^∞}, one observes that, for every section $\zeta \in \mu_{p^\infty}(R)$, since ζ is a p^{th}-power root of unity in R and p is nilpotent in R, we have that $(\zeta - 1) \in \mathrm{Nil}(R)$, and hence ζ coincides with the identity element of μ_{p^∞} modulo the nilpotent ideal $I := (\zeta - 1)R$. $\square$

Proposition 12.2. There is a canonical action of the formal group $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$ on the moduli space $\mathfrak{T}^{\mathrm{gr}}$ minus the boundary. Given $R \in \mathrm{Nilp}_{\mathcal{O}_E}$, $Z \in \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}(R)$ and $x \in \mathfrak{T}^{\mathrm{gr}}(R)$, corresponding to a $\mathcal{D}$-enriched abelian scheme $A(x)/R$, the p-divisible group of the abelian surface $A(Z \cdot x)$ is $E_Z + A(x)[p^\infty]$, where the $+$ operation is the one defined in Proposition 11.12, and E_Z is the Kummer extension constructed in Proposition 11.10.

Proof. The p-divisible group $G := A(x)[p^\infty]$, together with the $\mathcal{O}_B$-action and polarization it carries as well as the trivializations of its multiplicative and étale parts, is a section of the stack $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ over R. As a consequence, given an element $Z \in \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee)(R)$, we can form an extension $E_Z \in \mathrm{EXT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ thanks to Proposition 11.10, and construct a new p-divisible group $E_Z + G \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R)$ via the operation we have introduced in Proposition 11.12.

It follows from Proposition 11.12(a) that G' comes with a slope filtration whose slope pieces are identified with those of G. Moreover, a nilpotent ideal I, we have that Z coincides with the identity element of $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$ (see Lemma 12.1), and E_Z consequently coincides with the trivial extension, so that G and G' get canonically identified thanks to Proposition 11.12(d): in other words, G' is a deformation of $G \in \mathrm{BT}_{\mathcal{O}_B,*}(R/I)$ over R. By Serre-Tate lifting theory, $A(x)/(R/I)$ deforms uniquely to a $\mathcal{D}$-enriched abelian surface $Z \cdot A$ over R whose p-divisbile group $(Z \cdot A)[p^\infty]$ is the deformation $E_Z + G$ of $G \in \mathrm{BT}(R/I)$. This surface, together with the structure carried by its p-divisible group G', is classified by a point $Z \cdot x : \mathrm{Spec}(R) \to \mathfrak{T}^{\mathrm{gr}}$. The fact that $x \mapsto Z \cdot x$ satisfies the axioms of a group action (associativity axiom and neutral element axiom) is ensured by Proposition 11.12(c) and Proposition 11.12(d). □

On the scheme $\mathfrak{T}^{\mathrm{gr}} \setminus D$, along with the action of μ_{p^∞} we have constructed in Proposition 12.2, there is an action by $\mathfrak{M}^{\mathrm{gr}} = \mathbb{Z}_p^\times \times \mathrm{GL}_1$ discussed in the previous section. Their interaction goes as follows.

Proposition 12.3. Given an R-point (a,b) of $\mathfrak{M}^{\mathrm{gr}} = \mathrm{Aut}(\mathcal{L}) \times \mathfrak{M}'$, an element $Z \in (\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty})(R)$, and a point $x \in \mathfrak{T}^{\mathrm{gr}}(R)$ away from the boundary, we have that $(a,b) \cdot Z \cdot (a^{-1}, b^{-1}) \cdot x = (a \cdot Z) \cdot x$, where $a \cdot Z := (\mathrm{Sym}^2_{\mathcal{O}_B}(a^{\vee,-1}) \otimes \mathrm{id})(Z)$.

Proof. The two abelian schemes $A(x)$ and $A((a^{-1}, b^{-1}) \cdot x)$ only differ because they carry different trivializations of the isoclinic pieces of their common μ-ordinary p-divisible group. When seen as elements as $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, we have that $A((a^{-1}, b^{-1}) \cdot x)[p^\infty] = a^\vee \cdot A(x)[p^\infty]$, where the dot stands for the $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$-action on $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ discussed in Proposition 11.8 (see also Remark 11.9).

If we now apply $E_Z + (-)$ to both sides, we get that $E_Z + A((a^{-1}, b^{-1})x)[p^\infty] = E_Z + a^\vee \cdot A(x)[p^\infty]$. This shows that the p-divisible group $A(Z \cdot (a^{-1}, b^{-1}) \cdot x))$ coincides with $E_Z + a^\vee \cdot A(x)[p^\infty]$ as an element of $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$. Recalling that the $+$ operation and the assignment $Z \mapsto E_Z$ are $\mathrm{End}_{\mathcal{O}_B}(\mathfrak{L}^\vee)$-equivariant (see Proposition 11.12), we conclude that $A((a,b) \cdot Z \cdot (a^{-1}, b^{-1}) \cdot x))$ has $(E_{a^{\vee,-1} \cdot Z}) + A(x)[p^\infty] \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$ as a p-divisible group.

We remark that, since the $E_Z + (-)$ operation leaves all slope pieces of the p-divisible group unvaried, and it is not sensitive to the trivialization of the connected-connected pieces of the p-divisible group, it is clear that the action of Z commutes with that of b, which is the reason why the latter does not appear in the commutation relation in the statement of the proposition. □

We are now ready to state the following result, regarding the interaction of the newly defined $(\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty})$-action with the weights of modular forms. To express it, we will make use of the interpretation of the former as an action by continuous function on the lattice $\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L})$, which has been thoroughly discussed in Subsection 11.4.4.2.

Theorem 12.4. Let R be some p-adically complete $\mathcal{O}_E$-algebra. Let $\lambda : \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \to \mathrm{GL}(V)$ and $\rho : \mathfrak{M}^{\mathrm{gr}} \to \mathrm{GL}(W)$ be representations of $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$ and $\mathfrak{M}^{\mathrm{gr}} = \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \times \mathfrak{M}'$ respectively, being V and W some R-vector bundles. Let $f : \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}) \to V$ be an $\mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L})$-equivariant continuous function. The $(\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty})$-action on $\mathfrak{T}^{\mathrm{gr}}$ defined in Proposition 12.2 induces an operator $\vartheta^f : \omega^\rho \to \omega^{\rho \otimes \lambda}$ over $\mathfrak{Y}^{\mu-\mathrm{ord}} \hat{\otimes} R$.

Proof. This is an immediate application of Proposition 11.20. The compatibility condition expressed in Proposition 11.20(c) is guaranteed precisely by the Proposition 12.3 □

Remark 12.5. Theorem 12.4 only constructs the ϑ operator away from the boundary. If $x \in \mathfrak{T}^{\mathrm{gr}}$ is a geometric point belonging to the boundary, we have that $A(x)[p^\infty]$, i.e.

the p-torsion part of the semi-abelian scheme classified by x, is not a p-divisible group (only its connected part is). Hence, there is no clear way of deforming it by an extension of $\mathfrak{L} \otimes \mathbb{Q}_p/\mathbb{Z}_p$ by $\mathfrak{L}^\vee \otimes \mu_{p^\infty}$, and, should such a deformation be suitably definable, applying Serre-Tate lifting theory – which is another important ingredient of our construction – would also be non-straightforward. Because of these substantial difficulties, the present dissertation is not able to address the delicate problem of continuing the $(\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty})$-action, and hence the ϑ operator, to the boundary.

Among all ϑ^f operators introduced, the base case is represented by the one corresponding to $f = \mathrm{id}$.

Definition 12.6. Given the representation $\lambda : \mathrm{Aut}_{\mathcal{O}_B}(\mathfrak{L}) \to \mathrm{GL}_R(\mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L} \otimes_{\mathbb{Z}_p} R))$, we define the ϑ operator to be the operator $\vartheta^f : \omega^\rho \to \omega^{\rho+\lambda}$ corresponding to $f = \mathrm{id} : \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}) \to \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L} \otimes_{\mathbb{Z}_p} R)$.

CHAPTER 13

Theta operators beyond the μ-ordinary locus: the GSp_4 case

1. The GSp_4 setting

In this section, we present the construction of p-adic modular forms on p-rank ≥ 1 locus of the Siegel threefold, following, for almost all the material presented, Pilloni's paper [11].

1.1. The Siegel threefold. Let $G = \mathrm{GSp}_4$, and fix a prime p. Let $X \to \mathrm{Spec}(\mathbb{Z}_p)$ denote a toroidal compactification of the Siegel threefold of neat level $K \leq G$, where $K = K_p K^p$ with $K_p = \mathrm{GSp}_4(\mathbb{Z}_p)$. The subgroup K_ℓ equals the maximal compact subgroup $\mathrm{GSp}_4(\mathbb{Z}_\ell)$ for all but finitely many primes; let N be the product of those finitely many primes (so that no level structure is present away from N). The threefold X carries a canonical semi-abelian surface $A \to X$. We denote by D the boundary divisor, and we write $Y := X \smallsetminus D$ for the non-compactified threefold, i.e. the open subscheme of X over which A is an abelian surface (actually, $A_{|Y} \to Y$ is the *universal* principally polarized abelian surface of level K).

The sheaf of invariant differentials of A, which we denote ω, is a rank-2 vector bundle on X, i.e. a torsor for the group $M := \mathrm{GL}_2$, which should be thought of as the standard Levi subgroup of $G = \mathrm{GSp}_4$, modulo its center. For each choice of $(k_1, k_2) \in \mathbb{Z}^2$, $k_1 \geq k_2$, we define the automorphic vector bundle $\omega^{(k_1,k_2)} := \mathrm{Sym}^{k_1-k_2}(\omega) \otimes \det^{k_2}(\omega)$ corresponding to the irreducible representation of highest weight (k_1, k_2) of the group M (compare with Example 11.3).

1.2. Hecke algebra. The cohomology spaces $H^i(X, \omega^{(k_1,k_2)}) \otimes \mathbb{Q}_p$ of the automorphic vector bundles are finite-dimensional vector spaces carrying an action of the Hecke algebra $\mathcal{H} = C_c^0(\mathrm{GSp}_4(\mathbb{A}_f)//K)$, where C_c^0 denotes the continuous, compactly supported functions. At each prime $\ell \nmid N$, the Hecke algebra $\mathcal{H}_\ell$ is $\mathbb{Z}[T_{\ell,0}, T_{\ell,1}, T_{\ell,2}, T_{\ell,0}^{-1}]$, being $T_{\ell,0}, T_{\ell,1}$ and $T_{\ell,2}$ the characteristic functions of the double cosets $\mathrm{diag}(\ell, \ell, \ell, \ell)$, $\mathrm{diag}(\ell^2, \ell, \ell, 1)$, and $\mathrm{diag}(\ell, \ell, 1, 1)$. Each of these Hecke operators $T_{\ell,i}$ admits an interpretation as a correspondence from $X \otimes \mathbb{Q}_p$ to itself; more precisely,

(a) the correspondence inducing $T_{\ell,0}$ classifies the isogenies $A \to A'$ whose kernel is the subgroup $A[\ell]$;
(b) the correspondence inducing $T_{\ell,1}$ classifies the isogenies $A \to A'$ whose kernel is a totally isotropic subgroup of $A[\ell^2]$, whose intersection with $A[\ell]$ has degree ℓ^3.
(c) the correspondence inducing $T_{\ell,2}$ classifies the isogenies $A \to A'$ whose kernel is a totally isotropic subgroup of $A[\ell]$.

To summarize, the operator $T_{\ell,i}$ is associated with isogenies whose kernel is a totally isotropic subgroup of $A[\ell^{M(i)}]$, where $M(0) = M(1) = 2$, and $M(2) = 1$. For $\ell \neq p$, these correspondences are immediately seen to be also defined integrally. For $\ell = p$, the integral definition is subtler, and the integrality properties of the operators are discussed in [11, §7].

For later use, we also introduce here the Hecke polynomial $Q_\ell(X) \in \mathcal{H}_\ell[X]$, which is given by

$$Q_\ell(X) = 1 - T_{\ell,2}X + \ell(T_{\ell,1} + (\ell^2+1)T_{\ell,0})X^2 - \ell^3 T_{\ell,2}T_{\ell,0}X^3 + \ell^6 T_{\ell,0}^2 X^4.$$

Remark 13.1. The Hecke polynomial is homogeneous of degree 0, if one assigns degree $-M(i)$ to the symbol $T_{\ell,i}$ and degree 1 to the symbol X.

1.3. Klingen level. Let $\mathfrak{X} \to \mathrm{Spf}(\mathbb{Z}_p)$ be the formal completion of X at p, and let $\mathfrak{Y}$ be that of Y. We will work on the open formal loci $\mathfrak{X}^{\geq 1}$ and $\mathfrak{Y}^{\geq 1}$ where the p-rank of A is at least one.

The locus $\mathfrak{X}^{\geq 1}$ is stable under the action of Hecke operators; contary to the ordinary locus, it is not affine; instead, it is the union of two distinct, formal affine subschemes. This can be verified via the theory of partial Hasse invariants: the ordinary locus (i.e., the locus $\mathfrak{X}^{=2}$ where the p-rank is 2) is the complement, in $\mathfrak{X}$, of the zero-locus $\mathfrak{X}^{\leq 1}$ of the Hasse invariant $\mathrm{Ha} \in H^0(\mathfrak{X}\otimes\mathbb{F}_p, \det^{p-1}(\omega))$. Moreover, a second Hasse invariant $\mathrm{Ha}' \in H^0(\mathfrak{X}^{\leq 1}\otimes\mathbb{F}_p, \det^{p^2-1}(\omega))$ can be defined so that it vanishes exactly over the p-rank 0 locus: see [11, §6.3] for its definition. Now, since $\det(\omega)$ is an ample line bundle on $\mathfrak{X}\otimes\mathbb{F}_p$, a power $(\mathrm{Ha}')^k$ necessarily lifts (non-canonically) to a section $\widetilde{(\mathrm{Ha}')^k} \in H^0(\mathfrak{X}\otimes\mathbb{F}_p, \det^{k(p^2-1)}(\omega))$. Recalling that the complement of an ample divisor on a proper scheme is always affine, one can now conclude that $\mathfrak{X}^{\geq 1} = \{\mathrm{Ha} \neq 0\} \cup \{\widetilde{(\mathrm{Ha}')^k} \neq 0\}$ is the union of two affines.

For each $n \geq 1$, we can define a cover $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^n)} \to \mathfrak{X}^{\geq 1}$ classifying the subgroups $H_n \leq A[p^n]$ that are étale-locally isomorphic to μ_{p^n}. The map $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^n)} \to \mathfrak{X}^{\geq 1}$ is affine and étale (see [11, Lemma 9.1.1.1]), but it is not finite: over the ordinary locus $\mathfrak{X}^{=2} \subset \mathfrak{X}^{\geq 1}$, it has rank $p+1$, while it is an isomorphism when restricted to the closed subscheme $\mathfrak{X}^{=1} \subset \mathfrak{X}^{\geq 1}$ where the p-rank is 1. By taking the inverse limit, one obtains a pro-étale cover $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)} \to \mathfrak{X}^{\geq 1}$, which classifies all subgroups $H \leq A[p^\infty]$ that are pro-étale-locally isomorphic to μ_{p^∞}.

On $\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$, the inclusion of the universal multiplicative height-1 subgroup $H \hookrightarrow A[p^\infty]$ together with the dual projection $A[p^\infty] \twoheadrightarrow H^\vee$ induce a filtration $0 = \mathrm{Fil}_0(A[p^\infty]) \subset \mathrm{Fil}_1(A[p^\infty]) \subset \mathrm{Fil}_2(A[p^\infty]) \subset \mathrm{Fil}_3(A[p^\infty]) = A[p^\infty]$; the graded piece $\mathrm{gr}_2(A[p^\infty]) = H^\vee$ is étale of height 1, the graded piece $\mathrm{gr}_0(A[p^\infty]) = H$ is multiplicative of height 1, while the middle graded piece $\mathrm{gr}_1(A[p^\infty])$ is a p-divisible group of height 2 and dimension 1, which can alternatively be ordinary or connected-connected, depending on whether the p-rank of A is 1 or 2 (here, by $\mathrm{gr}_i(A[p^\infty])$ we mean $\mathrm{Fil}_{i+1}(A[p^\infty])/\mathrm{Fil}_i(A[p^\infty])$). At the level of differentials, we get an induced filtration $0 = \mathrm{Fil}^3(\omega) = \mathrm{Fil}^2(\omega) \subset \mathrm{Fil}^1(\omega) \subset \mathrm{Fil}^0(\omega) = \omega$, where $\mathrm{gr}^i(\omega) := \mathrm{Fil}^i(\omega)/\mathrm{Fil}^{i+1}(\omega)$ is the sheaf of invariant differentials of the p-divisible group $\mathrm{gr}_i(A[p^\infty])$. In particular, $\mathrm{gr}^2(\omega) = 0$, $\mathrm{gr}^1(\omega) = \ker(\omega \twoheadrightarrow \omega_H)$, $\mathrm{gr}^0(\omega) = \omega_H$. While the filtration $\mathrm{Fil}_i(A[p^\infty])$ only exists away from the boundary D, the filtration on the rank-2 vector bundle ω extends to the whole toroidal compactification $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}}(p^\infty)$.

If we let $M^{\mathrm{gr}} \leq M$ be the subgroup $\mathrm{GL}_1 \times \mathrm{GL}_1 \leq \mathrm{GL}_2$, over $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$ we have the left M-torsor $\mathcal{T}$ of the trivializations $\psi : \omega \xrightarrow{\sim} \mathcal{O}^2$, and the left M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$ of the trivializations $\psi_0 \oplus \psi_1 : \mathrm{gr}^0(\omega) \oplus \mathrm{gr}^1(\omega) \xrightarrow{\sim} \mathcal{O} \oplus \mathcal{O}$. For each choice of $(k_1, k_2) \in \mathbb{Z}^2$ with $k_1 \geq k_2$, we have

(1) a dominant weight (k_1, k_2) for M, corresponding to an irreducible representation V_{k_1,k_2} of M that, when twisted by the M-torsor $\mathcal{T}$ (seeSection 11.1), gives rise to the vector bundle $\omega^{(k_1,k_2)} = \mathrm{Sym}^{k_1-k_2}(\omega) \otimes \det^{k_2}(\omega)$ on $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$; and
(2) a character (k_1, k_2) for M^{gr}, i.e. a 1-dimensional representation $V^{\mathrm{gr}}_{k_1,k_2}$ of M^{gr}, which, once twisted by $\mathcal{T}^{\mathrm{gr}}$, gives rise to the line bundle $\mathrm{gr}^0(\omega)^{k_1-k_2} \otimes \det^{k_2}(\omega) \cong \mathrm{gr}^0(\omega)^{k_1} \otimes \mathrm{gr}^1(\omega)^{k_2}$, which we will denote by $\mathfrak{F}^{(k_1,k_2)}$.

The inclusion $M^{\mathrm{gr}} \leq M$ gives a surjective equivariant morphism $V_{(k_1,k_2)} \twoheadrightarrow V^{\mathrm{gr}}_{(k_1,k_2)}$, which in turn provides the canonical projection $\omega^{(k_1,k_2)} \twoheadrightarrow \mathfrak{F}^{(k_1,k_2)}$. This is an isomorphism only for scalar weights (i.e., when $k_1 = k_2$).

1.4. p-adic interpolation. Let now $\mathfrak{T}^{\mathrm{gr}}$ be the moduli space that classifies isomorphisms $\iota : \mu_{p^\infty} \xrightarrow{\sim} \mathrm{gr}_0 A[p^\infty]$ and $\psi_1 : \mathrm{gr}^1(\omega) \xrightarrow{\sim} \mathcal{O}$ over $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$: this is a $\mathfrak{M}^{\mathrm{gr}}$-torsor, where $\mathfrak{M}^{\mathrm{gr}} \leq M^{\mathrm{gr}}$ is the subgroup $\mathbb{Z}_p^\times \times \mathrm{GL}_1 \leq \mathrm{GL}_1 \times \mathrm{GL}_1$. An element $(a_0, a_1) \in \mathfrak{M}^{\mathrm{gr}}$ acts by taking $(\iota, \psi_1) \mapsto (\iota \circ a_0, a_1 \circ \psi_1)$. If we fix the standard trivialization $\mathcal{O} \cong \omega_{\mu_{p^\infty}}, 1 \mapsto dt/t$, then the differential of the map ι clearly induces an isomorphism $d\iota : \mathrm{gr}^0(\omega) \xrightarrow{\sim} \mathcal{O}$, and the assignement $(\iota, \psi_1) \mapsto (d\iota, \psi_1)$ consequently defines an equivariant morphism from the $\mathfrak{M}^{\mathrm{gr}}$-torsor $\mathfrak{T}^{\mathrm{gr}}$ to the M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$, which identifies $\mathcal{T}^{\mathrm{gr}}$ as the pushforward of $\mathfrak{T}^{\mathrm{gr}}$ along the inclusion $\mathfrak{M}^{\mathrm{gr}} \leq M$. In other words, the torsor $\mathfrak{T}^{\mathrm{gr}}$ we have just constructed is a $\mathfrak{M}^{\mathrm{gr}}$-reduction of the M^{gr}-torsor $\mathcal{T}^{\mathrm{gr}}$.

This reduction of the structure group allows for a (partial) p-adic variation of the weight (k_1, k_2). More precisely, given a continuous character $k_1 : \mathbb{Z}_p^\times \to R$ for R a p-adically complete $\mathbb{Z}_p$-algebra, and an integer $k_2 \in \mathbb{Z}$, we can form the continuous character $(k_1, k_2) : \mathfrak{M}^{\mathrm{gr}} \to R$, which gives rise to a line bundle $(\mathfrak{T}^{\mathrm{gr}})^{(k_1,k_2)}$ over $\mathfrak{X}^{\geq 1}_{\mathrm{Kli}(p^\infty)} \hat{\otimes}_{\mathbb{Z}_p} R$ (see Example 11.2), which we will still denote by $\mathfrak{F}^{(k_1,k_2)}$ (when $k_1 \in \mathbb{Z}$, this definition is compatible with the one given in the previous subsection). We will set $\mathfrak{F}^k := \mathfrak{F}^{(k,0)}$, so that $\mathfrak{F}^{(k_1,k_2)} = \mathfrak{F}^{k_1-k_2} \otimes \det(\omega)^{k_2} = \mathfrak{F}^{k_1} \otimes \mathrm{gr}^1(\omega)^{k_2}$. Finally, if we denote by $\kappa : \mathbb{Z}_p^\times \to \Lambda := \mathbb{Z}_p[[\mathbb{Z}_p^\times]]$ the universal character of $\mathbb{Z}_p^\times$, we have that any continuous p-adic character k of $\mathbb{Z}_p^\times$ is given by a continuous algebra homomorphism $k : \Lambda \to R$ with values in some p-adically complete $\mathbb{Z}_p$-algebra R, and $\mathfrak{F}^k = \mathfrak{F}^\kappa \hat{\otimes}_{\Lambda,k} R$. The sheaf $\mathfrak{F}^\kappa$ coincides with the one that Pilloni constructs in [11, §9.3] and denotes the same way.

2. The theta operator

We are now ready to construct a μ_{p^∞}-action on the torsor $\mathfrak{T}^{\mathrm{gr}}$.

Proposition 13.2. There is a canonical action of the formal group μ_{p^∞} on the moduli space $\mathfrak{T}^{\mathrm{gr}} \hat{\otimes} R$ minus the boundary. Given $x \in \mathfrak{T}^{\mathrm{gr}}(R)$, corresponding to an abelian surface $A(x)/R$, the p-divisible group of the abelian surface $A(\zeta \cdot x)$ is the p-divisible group $E_\zeta + A(x)[p^\infty]$, using the notation of Section 11.3.

PROOF. The proof is analogous to that of Proposition 12.2; we will consequently only briefly recall its main steps. The additional structure that $A[p^\infty]$ carries turns it into a section over R of the stack $\mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}$, with $\mathfrak{L} = \mathbb{Z}_p$. We can consequently form, thanks to

Proposition 11.10 and Proposition 11.12, a new p-divisible group $E_\zeta + A(x)[p^\infty]$, which is a deformation over R of $A(x)[p^\infty] \in \mathrm{BT}^{\mathfrak{L}}_{\mathcal{O}_B,*}(R/I)$, being I the nilpotent ideal $\zeta - 1$. This deformation of the p-divisible group of $A(x)$ induces a deformation of the whole abelian scheme thanks to Serre-Tate lifting theory: we set $\zeta \cdot x$ to be the point of $\mathfrak{T}^{\mathrm{gr}} \hat{\otimes} R$ classifying such deformation. $\square$

As we did in the μ-ordinary case, we now have to study the commutation relations between the newly introduced action and that of $\mathfrak{M}^{\mathrm{gr}}$.

Proposition 13.3. Given an R-point $(a,b) \in \mathfrak{M}^{\mathrm{gr}} = \mathbb{Z}_p^\times \times \mathrm{GL}_1$, an element $\zeta \in \mu_{p^\infty}(R)$, and a point $x \in \mathfrak{T}^{\mathrm{gr}}(R)$ away from the boundary, we have that $(a,b) \cdot \zeta \cdot (a^{-1}, b^{-1}) \cdot x = \zeta^{a^{-2}} \cdot x$.

Proof. The proof is analogous to that of Proposition 12.3, after we replace $Z \in \mathrm{Sym}^2_{\mathcal{O}_B}(\mathfrak{L}^\vee) \otimes \mu_{p^\infty}$ with $\zeta \in \mu_{p^\infty}$ (taking into account that $\mathfrak{L} = \mathbb{Z}_p$ in our setting). The fact that $\zeta^{a^{-2}}$ is the correct replacement for the expression $a \cdot Z$ appearing in the statement of Proposition 12.3 is explained in Remark 11.9. $\square$

Now that this commutation relation has been established, the ϑ operator is given as follows.

Theorem 13.4. Let $f : \mathbb{Z}_p \to R$ be a multiplicative function, and (k_1, k_2) be a pair with $k_1 : \mathbb{Z}_p^\times \to R$ a character (for R some p-adically complete $\mathbb{Z}_p$-algebra) and $k_2 \in \mathbb{Z}$. Then, the μ_{p^∞}-action defined in Proposition 13.2 defines an operator $\vartheta^f : \omega^{(k_1,k_2)} \to \omega^{(k_1+2f,k_2)}$ over $\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$.

Proof. This is an application of Proposition 11.20, using the following dictionary:

(a) $\mathfrak{S}$ is $\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}$;
(b) G is $\mathfrak{M}^{\mathrm{gr}} = \mathbb{Z}_p^\times \times \mathrm{GL}_1$;
(c) $\mathfrak{X}$ is $\mathfrak{T}^{\mathrm{gr}}$ minus the boundary;
(d) M is the lattice $\mathrm{Sym}^2(\mathbb{Z}_p) \cong \mathbb{Z}_p$, and G acts on it as follows: given $(a_0, a_1) \in \mathbb{Z}_p^\times \times \mathrm{GL}_1$, and $x \in \mathbb{Z}_p$, we have that $(a_0, a_1) \cdot x := a_0^{-2} x$;
(e) $V = R$, endowed with the G-action induced by the character $2f$ of $\mathbb{Z}_p^\times$;
(f) the G-equivariant function $f : M^\vee \to V$ is the multiplicative function $f : \mathbb{Z}_p \to R$ given in the statement of the theorem;
(g) W is the irreducible representation of highest weight (k_1, k_2) of G.

The compatibility condition expressed in Proposition 11.20(c) is guaranteed precisely by Proposition 13.3. $\square$

Remark 13.5. Among all ϑ^f operators introduced, the base case is represented by the one corresponding to $f = \mathrm{id}$, which gives an operator $\vartheta : \omega^{(k_1,k_2)} \to \omega^{(k_1+2,k_2)}$.

3. Interaction with Hecke operators

We are now going to discuss how the newly introduced ϑ operator affects Hecke eigenvalues.

To begin with, we introduce some notation. Let $R \in \mathrm{Nilp}_{\mathbb{Z}_p}$, let $x_1, x_2 \in \mathfrak{T}^{\mathrm{gr}}(R)$ be two points away from the boundary, and suppose we are given an isogeny $f : A(x_1) \to$

$A(x_2)$ between the abelian surfaces classified by those points. Assume that f respects the filtrations on $A(x_1)[p^\infty]$ and $A(x_2)[p^\infty]$, so that f restricts to a morphism $\mathrm{gr}_i f$ between the graded pieces $\mathrm{gr}_i(A(x_1)[p^\infty]) \to \mathrm{gr}_i(A(x_2)[p^\infty])$ for all $i = 0, 1, 2$. Both the domain and the codomain of $\mathrm{gr}_2 f$ are identified with $\mathbb{Q}_p/\mathbb{Z}_p$, hence $\mathrm{gr}_2 f$ is just the multiplication by some scalar $m_2(f) \in \underline{\mathbb{Z}_p}(R)$. Analogously, the domain and the codomain of $\mathrm{gr}_0 f$ are identified with μ_{p^∞}, hence $\mathrm{gr}_0 f$ is a power map $\mu_{p^\infty} \to \mu_{p^\infty}$, $\zeta \mapsto \zeta^{m_0(f)}$ for some $m_0(f) \in \underline{\mathbb{Z}_p}(R)$. Finally, since the sheaf of differentials of the p-divisible group $\mathrm{gr}_1(A(x_i)[p^\infty])$ is trivialized, the map $d(\mathrm{gr}_1(f)) : \mathrm{gr}_1(\omega_{x_2}) \to \mathrm{gr}_1(\omega_{x_1})$ is just $R \to R$, $a \mapsto m_1(f)a$ for some scalar $m_1(f) \in R$.

Lemma 13.6. The following hold.

(a) Let $x \in \mathfrak{T}^{\mathrm{gr}}(R)$ and $(a_0, a_1) \in \mathfrak{M}^{\mathrm{gr}}(R)$. Then, the identity morphism $f : A(x) \to A((a_0, a_1) \cdot x)$ satisfies $m_2(f) = a_0, m_1(f) = a_1^{-1}$ and $m_0(f) = a_0^{-1}$.

(b) Let $x_1, x_2, x_3 \in \mathfrak{M}^{\mathrm{gr}}(R)$, and let $f : A(x_1) \to A(x_2)$ and $g : A(x_2) \to A(x_3)$ be isogenies respecting the filtrations on p-divisible groups. Then, $m_i(gf) = m_i(g)m_i(f)$ for all $i = 1, 2, 3$.

(c) If $\lambda : A(x) \xrightarrow{\sim} A(x)^\vee$ is the principal polarization of which $A(x)$ is endowed, then $m_i(\lambda) = 1$ for all $i = 1, 2, 3$.

(d) If $N : A(x) \to A(x)$ is the multiplication-by-N isogeny, then $m_i(N) = N$ for all $i = 1, 2, 3$.

(e) Let $f : A(x_1) \to A(x_2)$ be an isogeny whose kernel is a totally isotropic subgroup of $A(x_1)[N]$, for some $N \geq 1$. Suppose that f respect the polarizations λ_i carried by $A(x_i)$, meaning that $f^\vee \lambda_2 f = N\lambda_1$. Then, $m_2(f)m_0(f) = N$

Proof. We omit the verification of (a), (b), (c), and (d). To prove (e), we notice that we have $f^\vee \lambda_2 f = N\lambda_1$ by assumption. If we apply m_0 to the left-hand side one obtains $m_0(f^\vee)m_0(\lambda_2)m_0(f) = m_2(f) \cdot 1 \cdot m_0(f)$. By applying it on the right-hand side, one gets $m_0(N)m_0(\lambda_2) = N \cdot 1$, whence the result follows. □

We are now ready to discuss how the μ_{p^∞}-action on $\mathfrak{T}^{\mathrm{gr}}$ interplays with isogenies between abelian surfaces.

Lemma 13.7. Given $R \in \mathrm{Nilp}_{\mathbb{Z}_p}$, take $\zeta_1, \zeta_2 \in \mu_{p^\infty}(R)$, and take $a, b \in \mathbb{Z}_p^\times$. Modulo the nilpotent ideal $I := (\zeta_1 - 1, \zeta_2 - 1)$, both Kummer extensions E_{ζ_1} and E_{ζ_2} are canonically identified with $\mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$. If $(\zeta_1)^a = (\zeta_2)^b$, the morphism $a \oplus b : \mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p \to \mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$ over R/I lifts to a (necessarily unique) morphism $g_{a,b} : E_{\zeta_1} \to E_{\zeta_2}$ over R. Such morphism restricts to a on μ_{p^∞}, and coincides with b on the quotient $\mathbb{Q}_p/\mathbb{Z}_p$.

Proof. Since E_{ζ_i} is the pushout of $0 \to \mathbb{Z}_p \to \mathbb{Q}_p \to \mathbb{Q}_p/\mathbb{Z}_p$ along $\zeta_i : \mathbb{Z}_p \to \mu_{p^\infty}$, and $\zeta_2 = \zeta_1^{a/b}$, by functoriality there exists a unique morphism $h : E_{\zeta_1} \to E_{\zeta_2}$ that restricts to ab^{-1} on the subgroup μ_{p^∞} and to the identity on the quotient $\mathbb{Q}_p/\mathbb{Z}_p$. If we compose h with the b-homothety morphism $b : E_{\zeta_2} \to E_{\zeta_2}$, we get the desired lift $g_{a,b}$. □

Proposition 13.8. Let $x_1, x_2 \in \mathfrak{T}^{\mathrm{gr}}(R)$, and suppose we are given an isogeny $f : A(x_1) \to A(x_2)$ between the abelian surfaces classified by those points, respecting the filtrations on the p-divisible groups $A(x_i)[p^\infty]$. Then, for all $\zeta \in \mu_{p^\infty}(R)$, we have that f deforms to a morphism $f' : A(\zeta^{m_2(f)} \cdot x_1) \to A(\zeta^{m_0(f)} \cdot x_2)$, i.e. there exists a unique such f' such that $f \equiv f' \mod (\zeta - 1)$. Moreover, f' is an isogeny that respects the filtrations on the p-divisible groups, and $m_i(f') = m_i(f)$ for all $i \in \{0, 1, 2\}$.

Proof. By Serre-Tate lifting theory it is enough to show that the deformation f' we are looking for can be built between the p-divisble groups of the two abelian schemes, i.e. that there exists $f' : A[p^\infty] + E_{\zeta^{m_2(f)}} \to A[p^\infty] + E_{\zeta^{m_0(f)}}$ such that $f \equiv f' \mod (\zeta - 1)$.

Now, let $g_{m_0(f), m_2(f)} : E_{\zeta^{m_2(f)}} \to E_{\zeta^{m_0(f)}}$ be the morphism constructed in Lemma 13.7. Since f and $g_{m_0(f),m_2(f)}$ have the same behavior once restricted to μ_{p^∞} and to $\mathbb{Q}_p/\mathbb{Z}_p$ (they both restrict to $m_0(f)$ and to $m_2(f)$, respectively), one can combine them to form a morphism $f' : A[p^\infty] + E_{\zeta^{m_2(f)}} \to A[p^\infty] + E_{\zeta^{m_0(f)}}$ between the Baer sum of their respective domains and codomains, which satisfies $m_i(f') = m_i(f)$ for all $i = 0, 1, 2$ (see Proposition 11.12(b)). Modulo $\zeta - 1$, $E_{\zeta^{m_0(f)}}$ and $E_{\zeta^{m_2(f)}}$ both coincide with the trivial extension $\mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$, $g_{m_0(f),m_2(f)}$ coincides with the morphism $m_0(f) \oplus m_2(f) : \mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p \to \mu_{p^\infty} \oplus \mathbb{Q}_p/\mathbb{Z}_p$ by construction, and f' consequently coincides with f. $\square$

Let us now fix a matrix $w = \mathrm{diag}(d_1, d_2, d_3, d_4) \in \mathrm{GSp}_4(\mathbb{Z})$, and assume all d_i are coprime with pN. We necessarily have that $d_3 d_2 = d_4 d_1 =: d(w)$. To such a coset, we can attach a correspondence $\mathfrak{C} \hookrightarrow \mathfrak{U} \times_{\mathbb{Z}_p} \mathfrak{U}$ that classifies the pairs of points $(x_1, x_2) \in \mathfrak{U} := \mathfrak{T}^{\mathrm{gr}} \smallsetminus D$ such that there exists a prime-to-p isogeny $f : A(x_1) \to A(x_2)$ of type w, meaning that

(a) the kernel $\ker(f)$ is a totally isotropic subgroup of $A(x_1)[d(w)]$, and the isogeny f respects the principal polarisations λ_i of the abelian surfaces $A(x_i)$, meaning that the following diagram commutes

$$\begin{tikzcd} A(x_1) \arrow[r, "f"] \arrow[d, "d(w)\lambda_1"] & A(x_2) \arrow[d, "\lambda_2"', "\cong"] \\ A(x_1)^\vee & A(x_2)^\vee \arrow[l, "f^\vee"] \end{tikzcd}$$

(b) for some appropriate choices of bases of the prime-to-p-Tate modules of $A(x_1)$ and $A(x_2)$, f is represented by the matrix w;

(c) the isogeny f respects the filtration on the p-disibible groups of $A(x_1)$ and $A(x_2)$, and its behaviour $m_i(f)$ on the trivialized graded piece $\mathrm{gr}_i(A[p^\infty])$ is fixed as follows: $m_0(f) = 1, m_1(f) = 1$ and $m_2(f) = d(w)$.

Let us further fix a weight (k_1, k_2), with $k_1 \in \mathbb{Z}_p^\times$, $k_2 \in \mathbb{Z}$. Let $\pi_1, \pi_2 : \mathfrak{C} \to \mathfrak{U}$ be the projections. The subscheme $\mathfrak{C} \hookrightarrow \mathfrak{U} \times_{\mathbb{Z}_p} \mathfrak{U}$ is invariant under the diagonal action of $\mathfrak{M}^{\mathrm{gr}}$, and that the maps $\pi_1, \pi_2 : \mathfrak{C} \to \mathfrak{U}$ are both $\mathfrak{M}^{\mathrm{gr}}$-equivariant. The correspondence induces a map in cohomology $T_w : R\Gamma(\mathfrak{U}; (k_1, k_2)) \to R\Gamma(\mathfrak{U}; (k_1, k_2))$, which is obtained by composing the pullback map $\pi_2^* : R\Gamma(\mathfrak{U}; (k_1, k_2)) \to R\Gamma(\mathfrak{C}; (k_1, k_2))$ and the trace map $\mathrm{tr}(\pi_1) : R\Gamma(\mathfrak{C}; (k_1, k_2)) \to R\Gamma(\mathfrak{U}; (k_1, k_2))$. This is the desired Hecke operator for the forms of weight (k_1, k_2).

Our aim is now studying the behavior of the correspondence T_w with respect to the action of μ_{p^∞}. Let us consider the action by μ_{p^∞} on $\mathfrak{U} \times_{\mathbb{Z}_p} \mathfrak{U}$ defined as follows: $\zeta \cdot (x_1, x_2) := (\zeta^{m_2(w)} \cdot x_1, \zeta^{m_0(w)} \cdot x_2)$. By Proposition 13.8, this action leaves the subscheme $\mathfrak{C}$ invariant. The interaction of the two projections π_1 and π_2 with the μ_{p^∞}-action is described in the following diagram:

$$\begin{array}{ccccc} \mathfrak{C} & \xrightarrow{\zeta} & \mathfrak{C} & \qquad \mathfrak{C} & \xrightarrow{\zeta} & \mathfrak{C} \\ \downarrow{\scriptstyle \pi_2} & & \downarrow{\scriptstyle \pi_2} & \qquad \downarrow{\scriptstyle \pi_1} & & \downarrow{\scriptstyle \pi_1} \\ \mathfrak{U} & \xrightarrow{\zeta^{m_0(w)}} & \mathfrak{U} & \qquad \mathfrak{U} & \xrightarrow{\zeta^{m_2(w)}} & \mathfrak{U} \end{array}$$

as a consequence, we get the following result at the level of cohomology:

$$\begin{array}{ccccc} R\Gamma(\mathfrak{U}) & \xrightarrow{\pi_2^*} & R\Gamma(\mathfrak{C}) & \xrightarrow{\mathrm{tr}(\pi_1)} & R\Gamma(\mathfrak{U}) \\ \downarrow{\scriptstyle \zeta^{m_0(w)}} & & \downarrow{\scriptstyle \zeta} & & \downarrow{\scriptstyle \zeta^{m_2(w)}} \\ R\Gamma(\mathfrak{U}) & \xrightarrow{\pi_2^*} & R\Gamma(\mathfrak{C}) & \xrightarrow{\mathrm{tr}(\pi_1)} & R\Gamma(\mathfrak{U}) \end{array}$$

In other words, we have that $T_w \circ \zeta = \zeta^{m_2(w)/m_0(w)} \circ T_w$. Meanwhile, $m_2(w)/m_0(w) = d(w)/1 = d(w)$. If we rephrase this commutation relation in terms of actions by $\mathrm{Cont}(\mathbb{Z}_p, R)$, we have proved the following proposition.

Proposition 13.9. Let $f : \mathbb{Z}_p \to R$ be a multiplicative continuous function, and let $\vartheta^f : \mathfrak{F}^{(k_1,k_2)} \to \mathfrak{F}^{(k_1+2f,k_2)}$ be the corresponding ϑ operator. Then, $T_w \vartheta^f = f(d(w))\vartheta^f T_w$.

This can in particular be applied to the operators $T_{\ell,0}$, $T_{\ell,1}$ and $T_{\ell,2}$ introduced in Subsection 13.1.1.2 to obtain commutation relations that are analogous to those presented in the introduction for the elliptic ϑ operator.

Corollary 13.10. The following commutation relations hold, for all primes $\ell \nmid Np$.

$$T_{\ell,0}\vartheta = \ell^2 \vartheta T_{\ell,0}$$
$$T_{\ell,1}\vartheta = \ell^2 \vartheta T_{\ell,1}$$
$$T_{\ell,2}\vartheta = \ell \vartheta T_{\ell,2}$$

4. Application to Galois representations

For each dominant weight $(k_1, k_2) \in \mathbb{Z}^2$, $k_1 \geq k_2$, one can compute the coherent cohomology groups $H^i(X, \omega^{(k_1,k_2)}) \otimes \mathbb{Q}_p$ and $H^i(X, \omega^{(k_1,k_2)}(-D)) \otimes \mathbb{Q}_p$ of the Siegel threefold, that are finite-dimensional $\mathbb{Q}_p$-vector spaces, endowed with an action of the Hecke algebra $\mathcal{H}$. Suppose now f is a cohomology class belonging to one of these cohomology groups, and that it is a simultaneous Hecke eigenform for all Hecke operators; its eigenvalues can be packed together into a homomorphism $\Theta_f : \mathcal{H} \twoheadrightarrow \overline{\mathbb{Q}}_p$.

Let now $\rho : G_{\mathbb{Q}} \to \mathrm{GSp}_4(E)$ be a p-adic Galois representation (being E some finite extension of $\mathbb{Q}_p$) that is unramified away from the primes dividing Np, and let $\sigma_\ell \in G_{\mathbb{Q}}$ be a Frobenius element relative to some prime $\ell \nmid Np$, i.e. σ_ℓ belongs to the decomposition group $D_{\mathfrak{L}}$ for some place $\mathfrak{L}$ of $\overline{\mathbb{Q}}$ above ℓ, and induces the Frobenius automorphism $x \mapsto x^\ell$ on the residue field $k(\mathfrak{L}) = \overline{\mathbb{F}_\ell}$. By unramifiedness, the conjugacy class of $\rho(\sigma_\ell)$ only depends on ℓ; in particular, for each prime ℓ one can compute the characteristic polynomial $\chi_\ell(\rho)(x) := \det(1 - x\rho(\sigma_\ell))$, which is a degree-4 polynomial with coefficients in $\overline{\mathbb{Q}}_p$ and constant term equal to 1.

Definition 13.11. We say that a semisimple continuous Galois representation $\rho : G_{\mathbb{Q}} \to \mathrm{GSp}_4(\overline{\mathbb{Q}}_p)$ unramified away from the primes dividing Np *is attached* to a Hecke eigensystem $\mathcal{H}^{Np} \twoheadrightarrow \overline{\mathbb{Q}}_p$ if $\chi_\ell(\rho) = \Theta(Q_\ell)$ for all $\ell \nmid Np$, being Q_ℓ the Hecke polynomial as defined in Subsection 13.1.1.2.

For Hecke eigensystems contributing to the coherent cohomology of the Siegel threefold X, we have the following existence result, which is a rephrasing of [11, Theorem 5.3.1].

THEOREM 13.12. Let $\Theta : \mathcal{H} \to E \subseteq \overline{\mathbb{Q}}_p$ be a Hecke eigensystem such that $\Theta = \Theta_f$ for some eigenform $f \in H^i(X \otimes \overline{\mathbb{Q}}_p, \omega^{(k_1,k_2)})$ or $f \in H^i(X \otimes \overline{\mathbb{Q}}_p, \omega^{(k_1,k_2)}(-D))$. Then, there exists a semisimple Galois representation $\rho_f : G_{\mathbb{Q}} \to \mathrm{GSp}_4(E)$ attached to Θ, unramified away from Np.

4.1. Twisting by cyclotomic characters. The abelian group $\mu_{p^\infty}(\overline{\mathbb{Q}})$ of p-power roots of unity in $\overline{\mathbb{Q}}$ is a $\mathbb{Z}_p$-module that is (non-canonically) isomorphic to $\mathbb{Q}_p/\mathbb{Z}_p$; its automorphism group is $\mathrm{GL}_1(\mathbb{Z}_p)$. The absolute Galois group $G_{\mathbb{Q}}$ clearly acts on $\mu_{p^\infty}(\overline{\mathbb{Q}})$, giving rise to a 1-dimensional Galois representation $\omega : G_{\mathbb{Q}} \to \mathrm{GL}_1(\mathbb{Z}_p) \hookrightarrow \mathrm{GL}_1(\mathbb{Q}_p)$, which is known as the *cyclotomic character*. Its main properties are presented in the following proposition.

Proposition 13.13. The cyclotomic character $\omega : G_{\mathbb{Q}} \to \mathrm{GL}_1(\mathbb{Q}_p)$ is unramified away from p. Moreover, if $\ell \neq p$ is another prime, and σ_ℓ is an ℓ-Frobenius element of $G_{\mathbb{Q}}$, we have that $\omega(\sigma_\ell) = \ell \in \mathrm{GL}_1(\mathbb{Q}_p)$.

PROOF. In $\mathbb{Q}_\ell^{\mathrm{ur}}$, the polynomial $x^{p^n} - 1$ has p^n distinct roots, and the Frobenius automorphism permutes them acting as $x \mapsto x^\ell$: this follows from Hensel's lemma. We deduce that $\mathbb{Q}(\mu_{p^\infty})$ is unramified at ℓ, and that the Frobenius element σ_ℓ acts as $x \mapsto x^\ell$ on p-th power roots. $\square$

Given a Galois representation $f : G_{\mathbb{Q}} \to \mathrm{GSp}_4(\overline{\mathbb{Q}}_p)$, one can compute its cyclotomic twist $f \otimes \omega$, which is defined as $(f \otimes \omega)(x) := \omega(x) f(x)$. We have the following properties.

Proposition 13.14. Let f and $f \otimes \omega$ be as above. Then, given a prime $\ell \neq p$, $f \otimes \omega$ is unramified at ℓ if and only if f is; moreover, in this case, $\chi_\ell(\rho \otimes \omega)(x) = \chi_\ell(\rho)(\ell x)$.

We are now ready to state the following theorem about the relation between applying the theta operator ϑ we have defined, and cyclotomic twists of the attached Galois representations.

THEOREM 13.15. Let $f \in H^i(X, \omega^{(k_1,k_2)}) \otimes \overline{\mathbb{Q}}_p$ be an eigenform for the Hecke algebra $\mathcal{H}$, being $i = 0$ or $i = 1$; let ρ_f be its attached Galois representation (see Theorem 13.12). Let $g := \vartheta(f) \in H^i(\mathfrak{Y}^{\geq 1}_{\mathrm{Kli}(p^\infty)}, \mathfrak{F}^{(k_1+2,k_2)}) \otimes \overline{\mathbb{Q}}_p$. Assume $g \neq 0$. Then, g is still an eigenform for the prime-to-p Hecke algebra $\mathcal{H}^p$; moreover, there exists a Galois representation ρ_g attached to its Hecke eigensystem, which coincides with the cyclotomic twist of f (in other words, $\rho_g \cong \rho_f \otimes \omega$).

PROOF. It follows from the commutation properties between Hecke operators and ϑ (Corollary 13.10) that, if $a_{\ell,i}$ (being $i = 0, 1, 2$) is the Hecke eigenvalue of f with respect to $T_{\ell,i}$, then g is an eigenform for $T_{\ell,i}$ with respect to the eigenvalue $\ell^{M(i)} a_{\ell_i}$, being $M(0) = M(1) = 2$, and $M(2) = 1$. Let us denote by Θ_f and Θ_g the Hecke eigensystems of f and g respectively: recalling the homogeneity properties of the Hecke polynomial Q_ℓ described in Remark 13.1, we consequently have that $\Theta_g(Q_\ell)(x) = \Theta_f(Q_\ell)(\ell x)$.

The representation ρ_f is unramified at all $\ell \nmid Np$, and the characteristic polynomials of the Frobenii are $\chi_\ell(\rho_f)(x) = \Theta_f(Q_\ell)(x)$ (see Theorem 13.12). If we recall the description of the cyclotomic twist given in Proposition 13.14, we clearly have that $\rho_f \otimes \omega$ is also unramified at all $\ell \nmid Np$, and that $\chi_\ell(\rho_f \otimes \omega)(x) = \Theta_f(Q_\ell)(\ell x)$.

Putting everything together, we conclude that $\chi_\ell(\rho_f \otimes \omega)(x) = \Theta_g(Q_\ell)(x)$, which is to say that $\rho_f \otimes \omega$ is the Galois representation attached to the Hecke eigensystem of g. $\square$

Bibliography

[1] Ehud de Shalit and Eyal Z. Goren. “Theta operators on unitary Shimura varieties”. In: *Algebra & Number Theory* 13.8 (2019), pp. 1829–1877.

[2] Ellen Eischen. “p-adic Differential Operators on Automorphic Forms on Unitary Groups”. In: *Annales de l’institut Fourier* 62.1 (2012), pp. 177–243.

[3] Ellen Eischen and Elena Mantovan. “p-adic families of automorphic forms in the µ-ordinary setting”. In: *American Journal of Mathematics* 143 (2017), pp. 1–52.

[4] Ellen Eischen et al. “Differential operators and families of automorphic forms on unitary groups of arbitrary signature”. In: *Documenta Mathematica* 23 (2018), pp. 445–495.

[5] A. Grothendieck and M. Raynaud. “Modeles de neron et monodromie”. In: *Groupes de Monodromie en Géométrie Algébrique*. Springer Berlin Heidelberg, 1972, pp. 313–523.

[6] Sean Howe. “A unipotent circle action on p-adic modular forms”. In: *Transactions of the American Mathematical Society Series B* 7.6 (2020), pp. 186–226.

[7] Nicholas M. Katz. “A result on modular forms in characteristic p”. In: *Modular Functions of one Variable V*. Ed. by Jean-Pierre Serre and Don Bernard Zagier. Springer Berlin Heidelberg, 1977, pp. 53–61.

[8] Kai-Wen Lan. *Arithmetic Compactifications of PEL-Type Shimura Varieties*. Princeton University Press, 2013.

[9] Barry Mazur and Peter Swinnerton-Dyer. “Arithmetic of Weil curves”. In: *Inventiones mathematicae* 25 (1974), pp. 1–61.

[10] Ben Moonen. “Serre-Tate theory for moduli spaces of PEL type”. In: *Annales scientifiques de l’École Normale Supérieure* Ser. 4, 37.2 (2004), pp. 223–269.

[11] Vincent Pilloni. “Higher coherent cohomology and p-adic modular forms of singular weights”. In: *Duke Mathematical Journal* 169.9 (2020), pp. 1647–1807.