



JACOBS
UNIVERSITY

Explicit Second p -Descent on Elliptic Curves

by

Brendan Matthew Creutz

Submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy in Mathematics

Dissertation Committee:

Prof. Dr. Michael Stoll

Prof. Dr. Ivan Penkov

Dr. Tom Fisher

DECLARATION

I, Brendan Matthew Creutz, hereby confirm that I completed this thesis independently, that I have not heretofore presented this thesis to another department or university, and that I have listed all references used, and have given credit to all additional sources of assistance.

Date, Place

Signature

Acknowledgements

I would like to thank:

my advisor Michael Stoll for suggesting this project and for his invaluable help along the way;

my parents for all of their support;

my brother for unwittingly inspiring me to study mathematics; and

Nika: Mano meilė tau kaip didelis kalnas, neregėtų aukštumų siekiantis virš debesų.

Contents

Introduction	9
1. Organization	11
2. Notation	13
Chapter I. Motivation and Background	17
1. What is descent good for?	17
2. The Selmer group	20
3. Using functions on the curve	23
4. How to do a p -descent on an elliptic curve	29
5. The interpretation as n -coverings	35
6. Projective models	40
Chapter II. The Descent Map	45
1. The fake Selmer set	46
2. The linear part of the descent map	48
3. The descent map	57
4. Injectivity of the descent map	60
5. Image of the descent map	65
6. The main diagram	69
7. Inverse of the descent map	71
Chapter III. Computing the p -Selmer Set	85
1. The algebraic Selmer set	85
2. The local image	89
3. Algorithms	91
4. Examples	94
5. Directions for further work	104
Bibliography	107

Introduction

One of the fundamental motivating problems in arithmetic geometry is to understand the set $V(k)$ of rational points on an algebraic variety V defined over a number field k . When $V = E$ is an elliptic curve, this set has a natural structure as a finitely generated abelian group (the Mordell-Weil group). The problem then becomes how to determine it in practice.

In addition to the Mordell-Weil group, there is another important arithmetic invariant of an elliptic curve: the Shafarevich-Tate group $\text{III}(E/k)$. An n -descent on an elliptic curve is a way to obtain information on both of these groups. For each integer $n \geq 2$, there is an exact sequence relating the two:

$$0 \rightarrow E(k)/nE(k) \rightarrow \text{Sel}^{(n)}(E/k) \rightarrow \text{III}(E/k)[n] \rightarrow 0.$$

The middle term is a finite group known as the n -Selmer group. An explicit n -descent on an elliptic curve computes the n -Selmer group and produces explicit representatives for its elements as curves in projective space. Determination of $\text{Sel}^{(n)}(E/k)$ yields partial information on the Mordell-Weil and Shafarevich-Tate groups. In addition, the models produced can often be used to find points of large height in the Mordell-Weil group or to study explicit counter-examples to the Hasse principle.

The main result of this thesis is an effective method for performing an explicit second p -descent on an elliptic curve when p is a prime. We assume an algorithm which performs an explicit p -descent on E , yielding models for the elements of $\text{Sel}^{(p)}(E/k)$ as genus one normal curves of degree p in \mathbb{P}^{p-1} . We then perform an explicit p -descent on some curve C thus obtained. This is a computation of the set $\text{Sel}^{(p)}(C/k)$ of everywhere locally solvable p -coverings of C which produces explicit models for its elements as genus one normal curves of degree p^2 in \mathbb{P}^{p^2-1} . Performing this computation for each element in $\text{Sel}^{(p)}(E/k)$ one obtains information that is just as good as that obtained by an explicit p^2 -descent on E .

As is typical for descents, the running time of our algorithm is dominated by the computation of class and unit group information in a certain étale k -algebra. In our situation this is the étale k -algebra of degree p^2 corresponding to the set of flex points of C . In addition, our algorithm requires computations in a second étale algebra of degree at most $p^2(p^2 - 1)/2$. The most expensive operation required there is the extraction of p -th roots of elements known to be p -th powers. When $p = 2$, this second algebra is simply k . When $p = 3$, one can get away with an algebra of degree 12, where such computations are entirely feasible. For larger p , however, this provides another barrier to the practical applicability of the algorithm.

The technique of descent to study solutions of Diophantine equations goes back at least to Fermat. In one of the first applications of computers to number theory Birch and Swinnerton-Dyer [BSD-I, -II] studied the Mordell-Weil groups of elliptic curves over \mathbb{Q} using 2-descents. These computations produced empirical evidence motivating their famous conjecture. Their method uses an explicit enumeration of certain homogeneous

spaces of the elliptic curve. While applicable (in principle) to larger n and over arbitrary number fields, the method is quickly defeated by combinatorial explosion when one ventures much beyond 2-descents over \mathbb{Q} .

There is an alternative approach which is based more closely on the original proof of the Mordell-Weil theorem [Mor, Weil]. First one computes the n -Selmer group as a finite subgroup of a finite exponent quotient of the multiplicative group of some étale k -algebra. One is then left with the task of constructing explicit models from the algebraic representatives. It is only in the past two decades or so that improved computing power, higher-level computer algebra software and better theoretical understanding have made computations using this alternative approach feasible. The first step requires deep arithmetic knowledge, such as S -class and μ -unit group information, of the constituent fields of the algebra. The most efficient known algorithm for obtaining this information in a number field has a running time which is exponential in the degree of the field.

Typically the algebra is related in some way to the group $E[n]$ of n -torsion points on E . For arbitrary n there is an algorithm involving the étale algebra $R = \text{Map}_k(E[n] \times E[n], \bar{k})$ of Galois equivariant maps from $E[n] \times E[n]$ to an algebraic closure of k [CFOSS-I, 3.2]. Typically R contains an extension of k of degree $O(n^4)$ making the arithmetic computations infeasible in practice. In general one can reduce computation of the n -Selmer group to the case that n is a prime power. For $n = p$ a prime, there is a method using the étale k -algebra $A = \text{Map}_k(E[p], \bar{k})$. Generically this splits as a product of k with some field extension of degree $p^2 - 1$. The p -Selmer group is then computed as a subgroup of $A^\times/k^\times A^{\times p}$. The situation for $n = 2$ is described in [Sim] and in [Sch2, St1] where 2-descent on Jacobians of hyperelliptic curves is also considered. For odd p , this was developed in the papers [DSS, SchSt]. For $p = 2, 3$, these algorithms are practical over number fields of moderate degree and discriminant and are part of the MAGMA computer algebra package.

For larger p these computations may still be feasible if favorable circumstances prevail. In particular, when E is p -isogenous to some other elliptic curve E' , the étale algebra in question may split making the computations much easier. Alternatively, one can combine the information from p -isogeny descents on E and E' . Among others, we mention here the works [Ba, CP, Fi1, Fi2, FG, Go, Sel, Ste, Top].

In the second step, one starts with representatives for the n -Selmer group in some étale algebra and wants to construct explicit models for the corresponding coverings. For $n \geq 3$, the problem is studied in the series of papers [CFOSS-I, -II, -III], the situation for $n = 2$ having been well-known for some time [Ca4, Section 15]. Starting with an element of R^\times , the multiplicative group of the étale k -algebra associated to $E[n] \times E[n]$, representing an element of the n -Selmer group, they show how to compute a collection of homogeneous equations defining a model for the corresponding covering as a genus one normal curve of degree n in \mathbb{P}^{n-1} . When n is prime it is also shown how the representatives for $\text{Sel}^{(n)}(E/k)$ in A^\times computed by the method mentioned above can be converted to representatives in R^\times . Taken together, this gives a complete method for performing explicit p -descents on elliptic curves.

The biggest obstacle to the practical implementation of their method is the need to find a rational point on an everywhere locally solvable Brauer-Severi variety of dimension $n - 1$. For $n = 2$, this means finding a point on a conic. In general the problem is equivalent to finding an explicit trivialization of a central simple k -algebra known to be k -isomorphic to a matrix algebra. For $n = 3$ and $k = \mathbb{Q}$, the authors of [CFOSS-I] have developed a practical method which is part of the 3-descent implementation in

MAGMA.

To our knowledge, the only existing practical methods for computing the n -Selmer group of a general elliptic curve when n is a higher prime power are for $n = 4$ [**MSS**, **Wom**] and $n = 8$ [**Sta**]. Rather than performing a direct 2^m -descent, these proceed by performing 2-descents in a tower of coverings. For example, the output of an explicit 2-descent on E is a finite collection of double covers of \mathbb{P}^1 ramified in four points. A second 2-descent computes the collection of everywhere locally solvable 2-coverings of one (or all) of these. The running time is dominated by the computation of arithmetic information in the étale algebra corresponding to the ramification points of the double cover of \mathbb{P}^1 . This is typically a field of degree 4, making the algorithm far more efficient than a direct 4-descent on E . The output of a second 2-descent is a finite collection of quadric intersections in \mathbb{P}^3 . These become the input for Stamminger's method for third 2-descent.

The method for second 2-descent described in this thesis is a (very) slight modification of the method above. While no more efficient, it does admit a cleaner cohomological interpretation which better integrates with our method for second p -descents for odd p . The primary reason for including it here is in the hope that the reader already familiar with second 2-descents will find the presentation for odd p more easily digestible. It does not appear that higher p -descents for odd p have been treated before in any systematic way.

Our algorithm is practical for $p = 3$ and $k = \mathbb{Q}$ and has been implemented by the author in MAGMA. Using this we are now able to exhibit explicit examples of elements of order 9 in the Shafarevich-Tate group of an elliptic curve. Alternatively, if the 3-primary part of the Shafarevich-Tate group has exponent 3, then the algorithm can be used to prove this unconditionally. For larger p the algorithm is decidedly less efficient although still approachable for $p = 5$. In an example we use a second 5-descent to prove that $\text{III}(E/\mathbb{Q})[5^\infty] = \text{III}(E/\mathbb{Q})[5] \neq 0$ for an elliptic curve E which admits a 5-isogeny over \mathbb{Q} . Combining our second 3-descent with the existing algorithms for higher 2-descents and some very deep results of Coates, Rubin, Wiles, et. al. we verify the full Birch and Swinnerton-Dyer conjecture for an elliptic curve over \mathbb{Q} with Shafarevich-Tate group of order 144.

1. Organization

Chapter I. This chapter contains the necessary background information for our investigations. For the most part the material is, if not classical (e.g. to be found in Silverman's book [**Sil**]), then at least well-known to the experts. The only possible exceptions being proposition 3.1, which is more general than the typical formulations found in the literature, and some details appearing in sections 5 and 6.

In section 1 we briefly motivate the desire to perform descent on elliptic curves. In the following section the n -Selmer and Tate-Shafarevich groups of an elliptic curve over a number field are defined using Galois cohomology and the standard proof of finiteness of the n -Selmer group is given. The general philosophy of using a Galois equivariant family of functions on a curve to study its Picard group is espoused in section 3. To allow for greater flexibility in dealing with such families we introduce the notions of a derived G_K -set and its corresponding induced norm map. In the following section we review the method described in [**SchSt**] for computing the p -Selmer group of an elliptic curve.

In the remaining two sections, the geometric interpretation in terms of n -coverings is considered. Most of this is based in one way or another on the material in [CFOSS-I, Sections 1-3]. To set the stage for second descents, we have extended many of the notions appearing there for elliptic curves to the case of genus one curves. Section 5 consists mainly of a theoretical description of these objects, whereas in section 6 we focus on period-index issues and concrete realizations of these abstract objects as curves in projective space.

Chapter II. Here we develop the theoretical basis for the algorithm and the cohomological interpretation of second p -descents. For the most part we work with a fixed genus one normal curve C of degree p defined over an arbitrary perfect field of characteristic not equal to p . The main object of study is the set of isomorphism classes of p -coverings of C with trivial obstruction (in a sense analogous to that of [CFOSS-I]). The primary tool is the *descent map*, which gives a concrete algebraic realization of this rather abstractly defined set.

In section 1 we outline a naive attempt at second p -descent. For $p = 2$ this reduces to the method of second 2-descent described in [MSS, Wom, BS, Sta]. It allows one to compute the 2-Selmer set of C up to sign, which is enough to recover everything we are after. For $p \geq 3$ the ambiguity becomes more pronounced and this naive method is ultimately unsuccessful. This motivates the following sections, where this naive method is refined and the ambiguity is eliminated. In section 2 we identify the domain of the descent map as a principal homogeneous space for a certain subgroup of $H^1(K, E[p])$. We give both a cohomological description of this subgroup and explicit representations of its members as elements of the multiplicative group of a certain étale K -algebra H . The descent map is then defined in section 3 as a map taking values in a certain quotient of H^\times . Ultimately we will see that the descent map may be interpreted as an affine map (loosely speaking a linear map followed by a translation), so that the material of section 2 can be understood as a study of its ‘linear part’. In sections 4 and 5 we show that the descent map is injective and determine its image. All of this is tied together by the ‘main diagram’ presented in section 6. Then in section 7, we construct an explicit inverse to the descent map. In particular, we show how to obtain explicit models for elements in the image of the descent map as genus one normal curves of degree p^2 in \mathbb{P}^{p^2-1} .

Chapter III. Here we specialize to the case that the base field is a number field. Armed with the material of the preceding chapter, computation of the Selmer set is almost routine. The descent map gives a bijection from $\text{Sel}^{(p)}(C/k)$ onto its image, which we call the algebraic p -Selmer set. This gives an algebraic presentation of the p -Selmer set which is amenable to machine computation. As with many descent algorithms, the first step is to reduce computation of the finite set of coverings that are locally solvable at all primes outside a certain finite set of primes S to S -class and -unit group computations. One can then deal with each of the remaining primes individually. These two steps are the topics of sections 1 and 2. The complete algorithm is then outlined in section 3. We then conclude with a small selection of examples in section 4 and a short discussion of possibilities for future work in section 5.

2. Notation

If \mathcal{G} is an abelian group (written additively or otherwise) and $n \geq 1$, we use $\mathcal{G}[n]$ to denote the subgroup of \mathcal{G} consisting of elements which are killed by n . For a commutative ring R we use R^\times to denote the multiplicative group of invertible elements. For n indivisible by the characteristic of R , the n -torsion subgroup of R^\times will be denoted $\mu_n(R)$ rather than $R^\times[n]$.

The symbol K will always denote a perfect field and p will always denote a prime number not equal to the characteristic of K . We always assume we have a fixed algebraic closure \bar{K} of K and any algebraic extension of K we write down is taken to be a subfield of \bar{K} . We write G_K for the absolute Galois group of K . We usually abbreviate $\mu_n(\bar{K})$ to μ_n . Since we restrict to perfect fields, the term local field will be used to mean the completion of a number field at some prime.

In the special case that $K = k$ is a number field, we make the following additional conventions. For each non-archimedean prime v of k , we fix some extension w of v to \bar{k} . This amounts to choosing a decomposition group $G_v \subset G_k$ (via the rule $\sigma \in G_v \Leftrightarrow w^\sigma = w$). If \bar{k}_v denotes the union of the completions of the finite subextensions of \bar{k} with respect to w , then it is an algebraic closure of k_v with Galois group G_v . Note however that \bar{k}_v is not a completion of \bar{k} with respect to w . For a non-archimedean prime v we use k_v^{unr} to denote the maximal unramified extension of k_v . This is an infinite Galois extension with group isomorphic to $\hat{\mathbb{Z}}$, the profinite completion of the integers. The group $I_v = \text{Gal}(k_v|k_v^{\text{unr}}) \subset G_v$ is called the inertia group at v .

Galois Cohomology. In this thesis we make constant use of Galois cohomology. For definitions and basic results we refer the reader to [Ser3, Chapter II]. We always consider G_K as a profinite group with the profinite topology. A G_K -set is a discrete topological space with a continuous action of G_K . A G_K -module is a commutative group object in the category of G_K -sets, i.e. a discrete abelian group M with a continuous action of G_K acting by automorphisms. The i -th Galois cohomology group with coefficients in M will be denoted by $H^i(G_K, M)$ or simply $H^i(K, M)$. The group $H^0(K, M)$ is the subgroup of elements of M invariant under the action of G_K . We will often denote this also by M^{G_K} .

If L is any field extension of K , the base change $K \rightarrow L$ induces a canonical homomorphism $\text{res}_{L/K} : H^i(K, M) \rightarrow H^i(L, M)$, which we refer to as the restriction map. There is also an injective inflation map, $\text{inf}_{L/K} : H^1(\text{Gal}(L|K), M^{G_L}) \rightarrow H^1(K, M)$, whose image is the kernel of $\text{res}_{L/K}$. When $K = k$ is a number field and $L = k_v$ is some completion, the restriction map is given, upon identification of $H^i(k_v, M)$ with $H^i(G_v, M)$, by restricting a cocycle defined on G_k to the subgroup $G_v \subset G_k$. In particular, the Galois cohomology groups do not depend on the choice of decomposition group.

If k_v is the completion of k at some non-archimedean prime v and M is a G_{k_v} -module, the unramified subgroup of $H^1(k_v, M)$ is defined to be the kernel of the restriction map $\text{res}_{k_v^{\text{unr}}/k_v} : H^1(k_v, M) \rightarrow H^1(I_v, M)$. By exactness of the inflation-restriction sequence this is isomorphic to $H^1(\text{Gal}(k_v^{\text{unr}}|k_v), M^{I_v})$. If M is a G_k -module, we say that an element of $H^1(k, M)$ is unramified at v if its image under $\text{res}_{k_v/k}$ lands in the unramified subgroup. For any set of primes S containing all archimedean primes if the exponent of M is even¹, we use $H^1(k, M; S)$ to denote the subgroup of elements that are unramified at all primes not in S .

¹If the exponent of M is odd and v is an archimedean prime, then the group $H^i(k_v, M)$ is trivial.

Divisors. Let C be a smooth, projective and absolutely irreducible curve defined over K . For a commutative K -algebra A , we write $C \otimes_K A$ for the scheme $C \times_{\text{Spec}(K)} \text{Spec}(A)$. When $A = \bar{K}$, we also write $\bar{C} = C \otimes_K \bar{K}$. We use $\kappa(\bar{C})$ and $\kappa(C)$ to denote the function field of \bar{C} and its G_K -invariant subfield, respectively. We use $\text{Div}(\bar{C})$ to denote the free abelian group on the set of \bar{K} -points of C . Its elements are called divisors and will often be written as integral linear combinations of points. If we wish to make it clear that we are considering a point as a divisor, we will use square brackets. So $[P] \in \text{Div}(\bar{C})$ is the divisor corresponding to $P \in C(\bar{K})$. The action of G_K on points extends to an action on divisors. We use $\text{Div}(C)$ for the G_K -invariant subgroup and refer to its elements as K -rational divisors. A closed point of the K -scheme C corresponds to a Galois orbit of points in $C(\bar{K})$. As such a closed point may be interpreted as an element of $\text{Div}(C)$. In fact, $\text{Div}(C)$ is the free abelian group on such closed points. We denote the divisor of a function $f \in \kappa(\bar{C})^\times$ by $\text{div}(f)$.

Two divisors are said to be linearly equivalent if their difference is equal to $\text{div}(f)$ for some rational function $f \in \kappa(\bar{C})^\times$. The group of principal divisors is $\text{Princ}(\bar{C}) = \{\text{div}(f) : f \in \kappa(\bar{C})^\times\}$. It follows from Hilbert's Theorem 90 that the G_K -invariant subgroup, $\text{Princ}(C) = \text{Princ}(\bar{C})^{G_K}$, is the group of divisors that are divisors of functions in $\kappa(C)^\times$. We use $\text{Pic}(\bar{C})$ to denote the group of divisors modulo principal divisors. Its G_K -invariant subgroup is denoted $\text{Pic}_K(C)$. This is not, generally speaking, the same as the group $\text{Pic}(C) := \text{Div}(C)/\text{Princ}(C)$; not every K -rational divisor class can be represented by a K -rational divisor (for an example see [Ca3]). There is however an injective map $\text{Pic}(C) \rightarrow \text{Pic}_K(C)$. In most of our applications this map is also surjective, so we will often (but not always) assume this is the case.

Since the degree of the divisor associated to a rational function is 0, there is a well-defined notion of degree for divisor classes in $\text{Pic}(\bar{C})$. We denote the set of divisor classes of degree $i \in \mathbb{Z}$ by $\text{Pic}^i(\bar{C})$. We use similar notation for the other groups defined above. The group $\text{Pic}_K^0(C)$ may be identified with the group of K -rational points on the Jacobian of C .

Étale K -algebras and G_K -Sets. If $K \subset L$ is an extension of fields and A is a K -algebra, then $A \otimes_K L$ is an L -algebra which we will denote simply by A_L . In the particular case $L = \bar{K}$, the notation \bar{A} will also be used to denote $A \otimes_K \bar{K}$. If $\phi : A \rightarrow B$ is a morphism of K -algebras, the induced map $A_L \rightarrow B_L$ will also be denoted by ϕ .

If Ω is a finite G_K -set, define $\bar{A}(\Omega) = \text{Map}(\Omega, \bar{K})$ to be the \bar{K} -algebra of maps from Ω to \bar{K} . There is a natural action of G_K on $\bar{A}(\Omega)$ defined by

$$\phi^\sigma : x \mapsto \left(\phi(x^{\sigma^{-1}}) \right)^\sigma .$$

As a \bar{K} -algebra $\bar{A}(\Omega)$ is isomorphic to $\prod_{i=1}^{\#\Omega} \bar{K}$, but the action of G_K is twisted by the action on Ω . The G_K invariant subspace of $\bar{A}(\Omega)$ is the space of G_K -equivariant maps $\text{Map}_K(\Omega, \bar{K}) := \text{Map}(\Omega, \bar{K})^{G_K}$. This is an étale K -algebra; it splits as a product of finite extensions of K corresponding to the orbits in Ω . This defines an anti-equivalence between the categories of finite G_K -sets and étale K -algebras (see for example [Le1]). We can also recover $\text{Map}(\Omega, \bar{K})$ by tensoring with \bar{K} .

In this thesis we frequently find ourselves working with objects defined over such algebras, e.g. varieties, points, functions, etc. From a scheme-theoretic point of view this presents no difficulty. It will, however, be convenient to interpret these objects as Galois equivariant maps. For example suppose C is a K -variety and $A = \text{Map}_K(\Omega, \bar{K})$. Then $C \otimes_K A$ is a scheme over A . We can interpret a rational function $f \in \kappa(C \otimes_K A)^\times$ as a Galois equivariant map $\Omega \rightarrow \kappa(\bar{C})^\times$ or equivalently as a Galois equivariant family

of rational functions $f_\omega \in \kappa(\bar{C})^\times$ indexed by $\omega \in \Omega$. The Galois equivariance means that $(f_\omega)^\sigma = f_{\omega^\sigma}$ for all $\sigma \in G_K$. The divisor of f can be interpreted as a Galois equivariant map $\Omega \rightarrow \text{Div}(\bar{C})$, and so on.

We mention here the generalization of Hilbert's Theorem 90 to étale algebras which states that $H^1(K, \bar{A}^\times) = 0$. To prove it one uses Shapiro's lemma to reduce to the usual Theorem 90 for fields. Using this with the Kummer sequence (as one does for fields) one is lead to an isomorphism $H^1(K, \mu_n(\bar{A})) \simeq A^\times/A^{\times n}$.

Internal referencing. Certain items (e.g. theorems, sections, definitions, etc.) are numbered within each chapter. A reference to an item appearing in the same chapter will give this number. A reference to an item appearing in a different chapter will give in addition the chapter number as a Roman numeral. For example, Theorem 2.2 of chapter 1 will be referred to as 2.2 (in chapter I) or I.2.2 (in chapters II and III).

Motivation and Background

1. What is descent good for?

Let E be an elliptic curve over a number field k . The celebrated Mordell-Weil theorem tells us that $E(k)$ is a finitely generated abelian group, the Mordell-Weil group. This finite description opens up the tantalizing problem of making the proof effective: How can we compute a set of generators explicitly?

From the structure theory of finitely generated \mathbb{Z} -modules we have that

$$E(k) \simeq \mathbb{Z}^r \times E(k)_{\text{tors}},$$

where $r \geq 0$ is the rank of $E(k)$ and $E(k)_{\text{tors}}$ is the torsion part of $E(k)$. It follows from the theorem that $E(k)_{\text{tors}}$ is finite. Moreover, there are effective methods for computing it, which at least over \mathbb{Q} are usually quite efficient (see for example [Sil, VIII.7]). This reduces the problem of computing $E(k)$ to finding generators of the free part. To do this it is enough to determine the rank and find sufficiently many independent points of infinite order (see [Sik2]).

The naive strategy for doing this is to just look for points, checking along the way if they are of infinite order and independent from the ones already found. The problem is of course knowing when to stop; there is currently no proven method for determining the rank of an arbitrary elliptic curve over a number field. We should remark that the conjecture of Birch and Swinnerton-Dyer does suggest a solution to this problem and that parts of it have been proven for elliptic curves of tiny rank over \mathbb{Q} , and that this is already a tremendous achievement (e.g. [Maz, Theorem 3]).

If one is able to compute the rank, then this gives an effective procedure for computing a set of generators. Unfortunately even if one knows ‘when to stop searching’, this is hardly efficient. The difficulty stems from the fact that the generators may have extremely large height. So a naive search is unlikely to find them in a reasonable amount of time. An explicit n -descent on an elliptic curve is a computation that can be helpful in addressing both of these problems.

Bounding the rank. The proof of the Mordell-Weil theorem breaks into two steps. First one proves that for some (any) $n \geq 2$, the group $E(k)/nE(k)$ is finite. Then, using the theory of heights, one shows that this implies the finite generation of $E(k)$. Given a set of points which generate $E(k)/nE(k)$, this second step is effective (see [Sik2]). The first step, however, is not. The best effective result of the proof is the finiteness of the n -Selmer group. The proof will be reviewed in the next section where we also derive the exact sequence,

$$0 \rightarrow E(k)/nE(k) \rightarrow \text{Sel}^{(n)}(E/k) \rightarrow \text{III}(E/k)[n] \rightarrow 0.$$

From this we see that the n -Selmer group contains an isomorphic copy of $E(k)/nE(k)$. Our inability to determine the subgroup corresponding to $E(k)/nE(k)$ is related to the failure of the Hasse principle for genus one curves. Given a curve defined over k with

points over every completion of k there is, in general, no known effective procedure for deciding if the curve has a k -rational point.

In any event, the size of the n -Selmer group provides an upper bound for the Mordell-Weil rank. More precisely, the rank r of the Mordell Weil group satisfies

$$n^r = \frac{\#\mathrm{Sel}^{(n)}(E/k)}{\#\mathrm{III}(E/k)[n] \cdot \#(E(k)_{\mathrm{tors}}/nE(k)_{\mathrm{tors}})}$$

As mentioned above there are effective methods for determining the torsion subgroup, so $\#(E(k)_{\mathrm{tors}}/nE(k)_{\mathrm{tors}})$ is ‘known’. Thus the only obstacle to computing a sharp upper-bound is the n -torsion in the Shafarevich-Tate group. The aforementioned BSD conjecture predicts that $\mathrm{III}(E/k)$ is finite. So it should be possible to avoid the obstruction with a suitable choice of n . For this reason it is desirable to have efficient methods for performing n -descents for different values of n .

One can also compare the upper-bound obtained from an n -descent with the lower-bound obtained from a point search. If the two coincide then surely $\mathrm{III}(E/k)[n] = 0$. Alternatively, comparing the bounds obtained from multiple descents, one may be able to prove that $\mathrm{III}(E/k)[n] \neq 0$ for some n . The nontrivial elements here correspond to genus one curves which violate the Hasse principle. It is thus of interest to be able to produce explicit models for these curves.

Cassels constructed an alternating bilinear pairing on $\mathrm{III}(E/k)$ whose kernel is the divisible subgroup [Ca1, Ca2]. If $\mathrm{III}(E/k)$ is indeed finite, then this implies that its order is a perfect square [Sil, Exer. 10.20]. This allows for an interesting hypothetical scenario: Assume p is prime (for simplicity) and that lower- and upper-bounds r_{ps} and r_p are obtained by a point search and a p -descent, respectively. Then

$$\dim_{\mathbb{F}_p} \mathrm{III}(E/k)[p] \leq r_p - r_{\mathrm{ps}} .$$

Conjecturally the dimension here is even¹. If $r_p - r_{\mathrm{ps}}$ is odd, then confidence in BSD justifies letting your computer search a bit longer.

This discussion applies, for example, to the curves

$$E_q : y^2 = x^3 + qx .$$

If q is a prime congruent to 3, 5, 13, or 15 mod 16 one can show that the \mathbb{F}_2 -dimension of the 2-Selmer group is 2 [Sil, X.6.2]. One dimension is accounted for by the nontrivial 2-torsion point $(0, 0) \in E_q(\mathbb{Q})$ (there are no elements of order 4), so

$$\mathrm{rank}(E_q(\mathbb{Q})) + \dim_{\mathbb{F}_2} \mathrm{III}(E_q/\mathbb{Q})[2] = 1 .$$

Assuming finiteness, the rank must be one. These curves also illustrate how descent can be used in conjunction with the parts of BSD that have been proven. If the L -series of an elliptic curve over \mathbb{Q} vanishes at $s = 1$ to order $o \leq 1$, then the rank is equal to o and the Shafarevich-Tate group is finite (see, for example, [Maz, Theorem 3]). For all primes $q \equiv 3, 5, 13, 15 \pmod{16}$ and less than 100000, MAGMA will happily report that the analytic rank of $E_q(\mathbb{Q})$ is one². What this actually means is that the value of the L -series at $s = 1$ is very close to zero³, but that the value of the derivative

¹For this one also needs to know that the kernel of the pairing on the p -torsion subgroup consists precisely of those elements divisible by p . See [Fi3, Section 5]

²Presumably someone has (at least attempted to) come up with a proof that the L -series of the elliptic curves in this family have simple zeros at $s = 1$, but we were unable to find any references. Bremner and Cassels have found generators for the Mordell-Weil group of E_q for all primes $q \equiv 5 \pmod{8}$ less than 1000 [BC]

³Tom Fisher has pointed out that results on parity show that the value is in fact 0.

there is definitely nonzero. In any event, the order of vanishing is ≤ 1 for these curves, so $\text{III}(E_q/\mathbb{Q})$ is finite. Hence $\text{III}(E_q/\mathbb{Q})[2]$ has even \mathbb{F}_2 -dimension and so must be trivial, proving unconditionally that the rank is in fact one.

Finding points. The elements of the n -Selmer group may also be interpreted as unramified coverings of E . These coverings have the property that every k -rational point of E lifts to a k -rational point on exactly one of the coverings. If one is able to obtain models for these as curves in projective space, then the logarithmic height of a lifted point should be $O(n)$ times smaller than that of its image on E . This means that searching for points on the coverings should be more efficient than (and as effective as) searching for points on E directly. For this reason it is desirable to have efficient methods for performing descents for larger values of n .

To illustrate this consider the prime $68749 \equiv 13 \pmod{16}$ and the curve E_{68749} in the family discussed above. Using some very deep results, we have shown that the rank is one. This could also be achieved by finding a point of infinite order. A naive point search would eventually find the point⁴

$$P = \left(\frac{427723613901884041}{394673451632100}, \frac{-287804417946186514282008589}{7840736712769435119000} \right)$$

of infinite order, which together with $(0, 0)$ generates $E_{68749}(\mathbb{Q})$. We found it much quicker using the implementation of 4-descent (i.e. second 2-descent) in MAGMA. One of the coverings produced is the quadric intersection

$$C = \left\{ \begin{array}{l} 4z_1z_2 + z_2^2 + 4z_1z_3 + z_3^2 + 6z_1z_4 - 4z_2z_4 - 4z_3z_4 + 6z_4^2 = 0 \\ 6z_1^2 - 2z_1z_2 - z_2^2 + 2z_3^2 + 12z_1z_4 - 2z_2z_4 + 2z_3z_4 - z_4^2 = 0 \end{array} \right\} \subset \mathbb{P}^3,$$

together with a degree 16 map to E_{68749} . One quickly finds the point $(0 : 1 : 1 : 1)$ mapping to P .

REMARK: The algorithm developed in this thesis is explicit in that it produces models for these coverings when n is the square of a prime. We must admit however that we have not yet developed a theoretical basis for ensuring that the models produced are suitable for the task of finding large points. The error term involved in the height estimate above depends to a large extent on the model chosen. In order to get something useful in practice, one needs the coefficients of the defining equations to be small. Thanks to Tom Fisher and Michael Stoll we have some ad hoc methods which yield significant improvements in this direction, but there is still work to be done.

Making finite III effective. For genus one curves, determining solvability is equivalent to determining whether a curve is isomorphic to its Jacobian. One way of interpreting the conjectured finiteness of the Shafarevich-Tate group is the following:

*In any family of pair-wise everywhere locally isomorphic smooth projective curves defined over a number field k there are only finitely many k -isomorphism classes.*⁵

⁴If we were so motivated, we could find an example for which no point of infinite order would fit between the margins.

⁵This is known for curves of genus $\neq 1$. For genus 0, this amounts to saying that a quaternion algebra is determined by its ramification. For genus ≥ 2 it follows from the fact that there are only finitely many automorphisms (see [Maz, I.5]).

In addition to helping make the Mordell-Weil theorem explicit, descent may allow us to make this finiteness statement explicit as well. By way of example we offer the following theorem (taken more or less directly from [Maz, Theorem 1]), whose proof is admittedly much deeper than anything else appearing in this thesis.

THEOREM 1.1 (Rubin, Selmer). *Let C/\mathbb{Q} be the curve in \mathbb{P}^2 defined by the equation $3x^3 + 4y^3 + 5z^3 = 0$. If V/\mathbb{Q} is any variety such that $V \otimes \mathbb{Q}_p$ is \mathbb{Q}_p -isomorphic to $C \otimes \mathbb{Q}_p$ for every prime $p \leq \infty$, then V is \mathbb{Q} -isomorphic to exactly one of the 5 pairwise nonisomorphic curves:*

$$\begin{aligned} 3x^3 + 4y^3 + 5z^3 &= 0, \\ 2x^3 + 5y^3 + 6z^3 &= 0, \\ 2x^3 + 3y^3 + 10z^3 &= 0, \\ x^3 + 4y^3 + 15z^3 &= 0, \\ x^3 + y^3 + 60z^3 &= 0. \end{aligned}$$

Moreover, each of the curves in this list is everywhere locally isomorphic to C .

The last curve in the list has the obvious (and unique) rational point $(1 : -1 : 0)$ and can be identified with the Jacobian E of the other curves in the list. The theorem is equivalent to saying that $\text{III}(E/\mathbb{Q}) \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ and that these curves represent the $5 = \frac{9-1}{2} + 1$ elements of $\text{III}(E/\mathbb{Q})/\{\pm 1\}$. The hard part is showing that $\text{III}(E/\mathbb{Q})$ is finite. Since E has complex multiplication and its L -series does not vanish at $s = 1$ this follows from work of Coates and Wiles [CoWi]. Subsequent work of Rubin [Ru] allows one to conclude even that the exponent of $\text{III}(E/\mathbb{Q})$ is a power of 2 times a power of 3. The contribution of Selmer [Sel] (some decades earlier) was to perform the descent necessary to prove that $\text{III}(E/\mathbb{Q})[3^\infty] \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ and produce the curves appearing in the list. Somewhere along the way one must also check that $\text{III}(E/\mathbb{Q})[2] = 0$ (which it is).

2. The Selmer group

Let E be an elliptic curve over a number field k and $n \geq 2$. As mentioned above, there is an exact sequence

$$0 \rightarrow E(k)/nE(k) \rightarrow \text{Sel}^{(n)}(E/k) \rightarrow \text{III}(E/k)[n] \rightarrow 0.$$

In this section we define the terms of this sequence using Galois cohomology and outline the (effective) proof that $\text{Sel}^{(n)}(E/k)$ is finite.

The Kummer sequence. Let K be a perfect field with absolute Galois group G_K . For any $n \geq 2$ indivisible by the characteristic of K , one has a short exact sequence

$$0 \rightarrow \mu_n \rightarrow \bar{K}^\times \rightarrow \bar{K}^\times \rightarrow 0.$$

Galois cohomology associates to this a long exact sequence of cohomology groups. From this one obtains a second short exact sequence

$$0 \rightarrow K^\times/K^{\times n} \xrightarrow{\delta} H^1(K, \mu_n) \xrightarrow{n} H^1(K, \bar{K}^\times)[n] \rightarrow 0.$$

The map labeled δ is the connecting homomorphism defined as follows. For $a \in K^\times$, one chooses any n -th root $\alpha \in \bar{K}^\times$. The image of a is then the (class of the) coboundary

of α , which is a 1-cocycle taking values in μ_n . In symbols

$$\delta(a) : G_K \ni \sigma \mapsto \alpha^\sigma / \alpha \in \mu_n.$$

One can check that this does not depend on the choice of α . The famous ‘Theorem 90’ of Hilbert says that $H^1(K, \bar{K}^\times)$ is trivial. Hence we have an isomorphism $K^\times / K^{\times n} \simeq H^1(K, \mu_n)$. If $\mu_n \subset K^\times$, then one recovers the classical ‘Kummer theory’,

$$K^\times / K^{\times n} \simeq \text{Hom}(G_K, \mathbb{Z}/n\mathbb{Z}).$$

This says that the cyclic extensions of K of degree dividing n are in one to one correspondence with classes of elements of K^\times modulo n -th powers, the correspondence being given by adjoining n -th roots.

Now suppose E is an elliptic curve defined over K . Since multiplication by n defines a surjective homomorphism on the \bar{K} -points of E , we can form a ‘Kummer sequence’ for E ,

$$0 \rightarrow E(K)/nE(K) \xrightarrow{\delta} H^1(K, E[n]) \rightarrow H^1(K, E)[n] \rightarrow 0.$$

The connecting homomorphism is defined in a completely analogous way. To compute the image on a class of $E(K)/nE(K)$ represented by some point P , one chooses a lift of P to a point $Q \in E(\bar{K})$, such that $nQ = P$. Any two choices differ by an n -torsion point of E . So, for any $\sigma \in G_K$, the point $Q^\sigma - Q \in E(\bar{K})$ is n -torsion. One checks that the assignment $\sigma \mapsto Q^\sigma - Q$ is a cocycle whose class does not depend on the choices for P and Q . Unlike the \mathbb{G}_m case, the group $H^1(K, E)$ may be infinite.

The Selmer Group. We now specialize to the case that $K = k$ is a number field. As previously mentioned, the group $E(k)/nE(k)$ is finite. To prove this, one shows that it sits inside a finite subgroup of $H^1(k, E[n])$. This subgroup is defined by imposing local conditions.

Forming the Kummer sequence over k and all of its completions gives a commutative diagram with exact rows, where the vertical maps are given by restriction

$$\begin{array}{ccccccc} 0 & \longrightarrow & E(k)/nE(k) & \xrightarrow{\delta} & H^1(k, E[n]) & \longrightarrow & H^1(k, E)[n] \longrightarrow 0 \\ & & \downarrow & & \downarrow & \searrow \alpha & \downarrow \\ 0 & \longrightarrow & \prod_v E(k_v)/nE(k_v) & \xrightarrow{\prod_v \delta_v} & \prod_v H^1(k_v, E[n]) & \longrightarrow & \prod_v H^1(k_v, E)[n] \longrightarrow 0 \end{array}$$

We are interested in the image of $E(k)/nE(k)$. From the exactness of the rows, one sees that this sits inside the kernel of the diagonal map. This leads to the following definition.

DEFINITION 2.1. *The n -Selmer and Tate-Shafarevich groups of E are the groups*

$$\begin{aligned} \text{Sel}^{(n)}(E/k) &= \ker \left(H^1(k, E[n]) \xrightarrow{\alpha} \prod H^1(k_v, E)[n] \right) \quad \text{and} \\ \text{III}(E/k) &= \ker \left(H^1(k, E) \longrightarrow \prod H^1(k_v, E) \right). \end{aligned}$$

The exact sequence

$$0 \rightarrow E(k)/nE(k) \rightarrow \text{Sel}^{(n)}(E/k) \rightarrow \text{III}(E/k)[n] \rightarrow 0$$

follows immediately.

REMARK: For any separable isogeny $\phi : J' \rightarrow J$ of Abelian varieties over K , we can form a Kummer sequence as above. Over a number field, one may define the ϕ -Selmer group of J' in a completely analogous way. It sits in an exact sequence

$$0 \rightarrow J(K)/\phi(J'(k)) \rightarrow \text{Sel}^\phi(J'/k) \rightarrow \text{III}(J'/k)[\phi] \rightarrow 0.$$

The proof of the finiteness of $\text{Sel}^{(n)}(E/k)$ given below works, with the obvious modifications, to prove finiteness in this more general context as well.

We come to the fundamental result of this section.

THEOREM 2.2. *Let E be an elliptic curve over a number field k and $n \geq 2$, then $\text{Sel}^{(n)}(E/k)$ is finite and computable.*

To show that the n -Selmer group is finite, one first shows that it is contained in the unramified outside S subgroup, $H^1(k, E[n]; S)$, for an appropriate finite set of primes S . The criterion of Neron-Ogg-Shafarevich shows that one may take S to be the set of primes v where E has bad reduction or such that v is archimedean or lies above n . The theorem then follows from the general fact that if M is a finite G_k -module and S is a finite set of primes containing all archimedean primes, then $H^1(k, M; S)$ is finite (see [Ser3, II.6.2]).

For $M = E[n]$, the proof of this fact goes as follows. First one reduces to the case when the n -torsion of E is k -rational. If $k(E[n])$ is the n -division field of E (i.e. the smallest extension over which all n -torsion points are defined), then the group on the left in the inflation-restriction exact sequence,

$$H^1(\text{Gal}(k(E[n])|k), E[n]) \rightarrow H^1(k, E[n]) \rightarrow H^1(k(E[n]), E[n]),$$

is finite and, at least in principle, computable. It thus suffices to consider the group on the right

This means we can assume $k = k(E[n])$. Under this assumption, the action of G_k on $E[n]$ is trivial and so

$$H^1(k, E[n]) = \text{Hom}(G_k, E[n]) = \text{Hom}(G_k, \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}).$$

The group on the right corresponds to the collection of all Galois extensions of k with Galois group which may be embedded in $\mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$. The n -th roots of unity are contained in k (because the Weil pairing on $E[n]$ is G_k -equivariant), so these extensions can be described using Kummer theory. Each is obtained by adjoining the n -th roots of a pair of elements $a, b \in k^\times/k^{\times n}$. For $v \notin S$, such an extension will be unramified at v if and only if $\text{ord}_v(a) \equiv \text{ord}_v(b) \equiv 0 \pmod{n}$ (recall that S is assumed to contain all archimedean primes and all primes dividing n). Let

$$k(S, n) = \{a \in k^\times/k^{\times n} : \forall v \notin S, \text{ord}_v(a) \equiv 0 \pmod{n}\}.$$

Elements of $H^1(k, E[n])$ unramified outside of S correspond to extensions which are unramified outside S . So we have an isomorphism $H^1(k, E[n]; S) \simeq k(S, n) \times k(S, n)$. The effective proof that the $H^1(k, E[n]; S)$ is finite is completed by the following important theorem.

THEOREM 2.3. *Let k be a number field, $n \geq 2$ and S a finite set of primes containing all archimedean primes and all non-archimedean primes above n . Let*

$$k(S, n) = \{a \in k^\times/k^{\times n} : \forall v \notin S, \text{ord}_v(a) \equiv 0 \pmod{n}\}.$$

Then $k(S, n)$ is finite and there is an effective procedure for computing (a set of representatives in k^\times) for $k(S, n)$ which is efficient modulo computation of the n -torsion in the S -class group of k and the cokernel of multiplication by n on the S -unit group of k .

To prove this, one shows (see for example [St2], [PS]) that there is an exact sequence

$$0 \rightarrow \mathcal{O}_{k,S}^\times / \mathcal{O}_{k,S}^{\times n} \rightarrow k(S, n) \rightarrow \text{Cl}_S(k)[n] \rightarrow 0,$$

where $\mathcal{O}_{k,S}^\times$ and $\text{Cl}_S(k)$ denote the S -unit and S -class groups of k , respectively. Since the class group is finite and the S -unit group is finitely generated, all terms in the exact sequence are finite. The fact that this is effective follows from the fact that the S -class and -unit groups can be computed and that the maps in the short exact sequence above can be determined explicitly.

More generally, if A is an étale k -algebra, we will use the notation $A(S, n)$ to denote the unramified outside S part of $A^\times / A^{\times n}$. If A decomposes into a product of number fields $A \simeq \prod k_i$, then $A(S, n) \simeq \prod k_i(S, n)$. So the theorem applies in this situation as well.

According to Lenstra [Le2] there ‘appears to be’ a deterministic algorithm for determining the class and unit groups of a number field k in time at most $(2 + \log |\Delta|)^{O(d)} |\Delta|^{3/4}$, where $d = [k : \mathbb{Q}]$ is the degree of k and Δ is its discriminant. If one only needs the kernel and cokernel of multiplication by n , some improvement should be possible, but the consideration of S -class and -unit groups only makes things more complicated. In any event, the complexity should remain exponential in the degree.

The computation of $H^1(k, E[n]; S)$ outlined above requires one to extend to the n -division field of E . Generically this will be a $\text{GL}_2(\mathbb{Z}/n\mathbb{Z})$ extension of k . The large degree of this extension makes the resulting algorithm infeasible in most situations as computation of the S -class and -unit groups is not likely to be manageable. Even if it is, one still needs to somehow return to k , which may be problematic. Any practical algorithm should avoid this step.

In any event, once one has computed $H^1(k, E[n]; S)$ as a finite subgroup of $A(S, n)$, the n -Selmer group can be determined by considering the local conditions at the finitely many primes of S . The groups $H^1(k_v, E[n])$ are finite, so this is a finite problem to which there is an effective solution.

REMARK: This discussion remains valid in the more general context of computing the ϕ -Selmer group of an Abelian variety J . Namely, there is an effective procedure which requires S -class and -unit group information in the minimal extension of k over which all points in the kernel of ϕ and the relevant roots of unity are defined.

3. Using functions on the curve

In this section we show how one can study the Picard group of a curve by using functions on the curve. The K -rational points of the Jacobian of a curve C can be identified with $\text{Pic}_K^0(C)$, so one can use this to study the Mordell-Weil group. In the next section we will describe in detail how this leads to a far more efficient algorithm for computing the n -Selmer group of an elliptic curve when n is prime.

A general framework for doing this is described by Schaefer in [Sch2]. Though more general cases have been treated (notably cyclic covers of the projective line [PS]), at the time this encompassed virtually all known explicit methods for performing descent on Jacobians. Under certain assumptions, a rational function on C will induce a homomorphism from $J(K) = \text{Pic}_K^0(C)$ into a group of finite exponent. These induced maps are given quite explicitly by evaluating functions at points of C and as such are suitable for practical computations. Given an isogeny $\phi : J' \rightarrow J$, the connecting homomorphism in the corresponding Kummer sequence is also a homomorphism to a group of finite exponent. Schaefer presents a strategy for choosing functions so that the induced maps are related to (or can be identified with) a given connecting homomorphism.

In the context of [Sch2] the codomain of these induced maps is of the form $A^\times/A^{\times n}$, where A is some étale K -algebra and $J'[\phi] \subset J'[n]$. Schaefer shows that under certain assumptions one can choose the function(s) so that the induced map Φ factors as

$$\begin{array}{ccc} J(K) & \xrightarrow{\Phi} & A^\times/A^{\times n} \\ & \searrow \delta & \nearrow \beta \\ & H^1(K, J'[\phi]) & \end{array}$$

where δ is the connecting homomorphism and the map β is injective. Over a number field k , this ideal situation yields a commutative diagram

$$\begin{array}{ccc} J(k)/\phi J'(k) & \xrightarrow{\Phi} & A^\times/A^{\times n} \\ \downarrow & & \downarrow \\ \prod_v J(k_v)/\phi J'(k_v) & \xrightarrow{\prod \Phi_v} & \prod_v A_v^\times/A_v^{\times n} \end{array}$$

In principle, one can then compute $H^1(k, J'[\phi]; S)$ as a subgroup of $A(S, n)$. Having accomplished this, the explicit nature of the maps Φ_v makes computation of the ϕ -Selmer group (as a subgroup of $A(S, n)$) entirely practical.

We would like to build on this perspective in essentially two orthogonal directions. The first, which has been observed by Siksek [Sik3], is that under the assumptions of [Sch2] these functions actually induce maps on all of $\text{Pic}_K(C)$ and not just $\text{Pic}_K^0(C) = J(K)$. In the context of second p -descent this is useful because we are now interested in $\text{Pic}_K^1(C)$.

The second drives to the heart of the assumptions required for the method in [Sch2] to work, namely that the map labelled β (or equivalently Φ) in the diagram above be injective. When this is not satisfied information can get lost. Correcting for this requires allowing for other functions which induce homomorphisms from the Picard group to a seemingly more complicated group of finite exponent which is actually just as easy to work with in practice. Our description of this group involves an ‘induced norm map’ coming from the notion of a derived G_K -set which we now describe.

Derived G_K -sets. If Ψ, Ω are finite G_K -sets, we will say that Ψ is *derived from* Ω if the elements of Ψ are unordered tuples (multisets) of elements of Ω and the action on Ψ is induced by that on Ω . For example, the set of unordered pairs (distinct or not) of elements in Ω is a derived G_K -set. One can also interpret the elements of Ψ as formal integral linear combinations of elements of Ω with nonnegative coefficients. In this interpretation we write $b \in \Psi$ as a sum $\sum n_a a$ of distinct elements $a \in \Omega$.

Recall the notation of section 2 of the introduction. To Ω and Ψ we associate étale K -algebras, $A(\Omega) = \text{Map}_K(\Omega, \bar{K})$ and $A(\Psi) = \text{Map}_K(\Psi, \bar{K})$. If Ψ is derived from Ω as a G_K -set, then we define ‘induced norm maps’ between the corresponding algebras:

$$A(\Omega) = \text{Map}_K(\Omega, \bar{K}) \ni \phi \mapsto \left(b = \sum n_a a \mapsto \prod_a \phi(a)^{n_a} \right) \in \text{Map}_K(\Psi, \bar{K}) = A(\Psi),$$

$$A(\Psi) = \text{Map}_K(\Psi, \bar{K}) \ni \phi \mapsto \left(a' \mapsto \prod_{b=\sum n_a a} \phi(b)^{n_a} \right) \in \text{Map}_K(\Omega, \bar{K}) = A(\Omega).$$

These maps may also be interpreted as follows. Let Θ be the G_K -set whose elements are the integral multiples of elements of $\Omega \times \Psi$ of the form $n_a(a, b)$, where $b = \sum n_{a'} a'$ (the action on Θ is that derived from Ω). In terms of multisets, if a is in b with exact multiplicity n_a , then Θ contains the multiset which consists of (a, b) taken with multiplicity n_a . Now $A(\Theta)$ splits as a product $A(\Theta) = \prod_{\mathcal{O}} A(\mathcal{O})$ of finite extensions of K corresponding to the G_K -orbits $\mathcal{O} \subset \Theta$. If \mathcal{O} is some orbit, then all $n_a(a, b) \in \mathcal{O}$ have the same multiplicity (i.e. $n_a = n_{\mathcal{O}}$ is constant on \mathcal{O}). This is used to assign a weight $n_{\mathcal{O}}$ to each factor of $A(\Theta)$ and thus a map:

$$w : A(\Theta) = \prod_{\mathcal{O}} A(\mathcal{O}) \ni (\phi_{\mathcal{O}}) \mapsto (\phi_{\mathcal{O}}^{n_{\mathcal{O}}}) \in \prod_{\mathcal{O}} A(\mathcal{O}) = A(\Theta).$$

On the other hand, Θ comes equipped with projection maps

$$\Theta \ni n_a(a, b) \mapsto a \in \Omega$$

$$\Theta \ni n_a(a, b) \mapsto b \in \Psi.$$

These obviously respect the action of G_K , so one has corresponding inclusions of $A(\Omega)$ and $A(\Psi)$ into $A(\Theta)$. Thus $A(\Theta)$ comes naturally equipped with norm maps N_{Ω} and N_{Ψ} to $A(\Omega)$ and $A(\Psi)$, respectively. The induced norm maps above are given by the compositions:

$$A(\Omega) \hookrightarrow A(\Theta) \xrightarrow{w} A(\Theta) \xrightarrow{N_{\Psi}} A(\Psi),$$

$$A(\Psi) \hookrightarrow A(\Theta) \xrightarrow{w} A(\Theta) \xrightarrow{N_{\Omega}} A(\Omega).$$

While seemingly complicated, this description makes explicit computation of these induced norm maps in concrete examples rather straight-forward. Two examples are given following the proposition below; naturally, we will see more in the following chapter.

Descent on Picard groups using functions on curves. ⁶ Let C be a smooth, projective and absolutely irreducible curve defined over K . We want to use the notion of derived G_K -sets to study the group $\text{Pic}(C)$ of divisor classes on C which can be represented by a K -rational divisor.

Let $\Omega \subset C(\bar{K})$ be a finite G_K -set of geometric points on C and $\Psi \subset \text{Div}(\bar{C})$ a finite G_K -set of effective divisors on C which are supported on Ω . We use $A(\Omega)$ and $A(\Psi)$ to denote the corresponding étale K -algebras. As a G_K -set Ψ is derived from Ω and we have an induced norm map $\partial : A(\Omega) \rightarrow A(\Psi)$.

Now consider a rational function

$$f = (f_{\psi}) \in \kappa(C \otimes_K A(\Psi))^{\times} = \text{Map}_K(\Psi, \kappa(\bar{C})^{\times}).$$

⁶We have borrowed the name of this subsection from a short paper by Siksek [Sik3], where the case of a K -rational function is considered.

We consider this either as a function on $C \otimes_K A(\Psi)$ or as a Galois equivariant family of rational functions in $\kappa(\bar{C})^\times$ parameterized by $\psi \in \Psi$. We interpret the divisor of f , an element of $\text{Div}(C \otimes_K A(\Psi))$, as a G_K -equivariant map $\Psi \rightarrow \text{Div}(\bar{C})$. We write $\text{div}(f)$ as a difference of effective divisors. This can be interpreted as a difference of a pair of G_K -equivariant maps $[f]_0, [f]_\infty : \Psi \rightarrow \text{Div}(\bar{C})$ whose values at $\psi \in \Psi$ are the zero and pole divisors of f_ψ , respectively. Now suppose f satisfies

- (1) $\forall \psi \in \Psi, [f]_0(\psi) = \psi$, and
- (2) $\forall \psi \in \Psi$ and $\sigma \in G_K, [f]_\infty(\psi^\sigma) = [f]_\infty(\psi)$.

The first condition says that ψ is the zero divisor of the function $f_\psi \in \kappa(\bar{C})^\times$. The second condition amounts to saying that the map $[f]_\infty : \Psi \rightarrow \text{Div}(\bar{C})$ is constant on each G_K -orbit in Ψ . The Galois equivariance then implies that, on each orbit $\mathcal{O} \subset \Psi$, the value of $[f]_\infty$ is some K -rational divisor $d_{\mathcal{O}} \in \text{Div}(C)$. We write each as a sum

$$d_{\mathcal{O}} = \sum_{\mathcal{P}} m_{\mathcal{O}, \mathcal{P}} \mathcal{P},$$

of closed points \mathcal{P} and set $m_{\mathcal{O}} = \text{gcd}(m_{\mathcal{O}, \mathcal{P}})$ (recall that a closed point corresponds to a G_K -orbit of points in $C(\bar{K})$). Using these weights, we define a map

$$\iota : K \ni a \mapsto (a^{m_{\mathcal{O}}})_{\mathcal{O}} \in \prod_{\mathcal{O}} A(\mathcal{O}) = A(\Psi)$$

PROPOSITION 3.1. *With notation as above, let $d = \sum_P n_P [P] \in \text{Div}(C)$ (written as a sum of \bar{K} -points of C) be any K -rational divisor on C with support disjoint from all zeros and poles of the f_ψ .*

- (1) *Evaluating f on d gives a well-defined element $f(d) := \prod_P f(P)^{n_P} \in A(\Psi)^\times$.*
- (2) *f induces a unique homomorphism*

$$\text{Pic}(C) \rightarrow \frac{A(\Psi)^\times}{\iota(K^\times) \partial(A(\Omega)^\times)}$$

with the property that, for all d as above, the image of the class of d is equal to the class of $f(d)$.

Before giving the proof, it seems appropriate to provide a couple of familiar examples.

p -descent on elliptic curves. Let E be an elliptic curve over K and p a prime different from the characteristic of K . For each nontrivial p -torsion point P , one can find a function $f_P \in \kappa(\bar{C})^\times$ with $\text{div}(f_P) = p[P] - p[0_E]$. The existence of such a function is used in the definition of the Weil pairing ([Sil, III.8]). It follows from Hilbert's Theorem 90 that f can be defined over the field $K(P)$ obtained by adjoining the coordinates of P to K .

Let $A = \text{Map}_K(E[p] \setminus \{0_E\}, \bar{K})$ be the étale algebra corresponding to the nontrivial p -torsion points. Given a set $\{P_1, \dots, P_r\} \subset (E[p] \setminus \{0_E\})$ of representatives for the G_K -orbits and functions f_{P_i} as above, there is a unique way to extend the f_{P_i} to a Galois equivariant family of rational functions indexed by $E[p] \setminus \{0_E\}$. This gives a

function $f = (f_P) \in \kappa(C \otimes_K A)^\times$. In the notation of the proposition we have

$$\begin{aligned}\Omega &= E[p] \setminus \{0_E\}, \\ \Psi &= \{p[P] : P \in \Omega\}, \\ [f]_\infty &= p[0_E], \\ A(\Omega) &\simeq A(\Psi) \simeq A, \\ \partial &: \alpha \mapsto \alpha^p, \text{ and} \\ \iota &: a \mapsto a^p.\end{aligned}$$

Since $E(K) \neq \emptyset$, we have $\text{Pic}_K(E) = \text{Pic}(E)$. The proposition says that f induces a homomorphism

$$\text{Pic}_K(E) \rightarrow \frac{A^\times}{K^{\times p} A^{\times p}} = \frac{A^\times}{A^{\times p}}.$$

Since the target is of exponent p , this factors through the cokernel of multiplication by p on $\text{Pic}_K(E)$. So we get an induced homomorphism

$$\frac{E(K)}{pE(K)} = \frac{\text{Pic}_K^0(E)}{p\text{Pic}_K^0(E)} \hookrightarrow \frac{\text{Pic}_K(E)}{p\text{Pic}_K(E)} \rightarrow \frac{A^\times}{A^{\times p}}.$$

In section 4 we will see that the cohomology group $H^1(K, E[p])$ may be identified with a subgroup of $A^\times/A^{\times p}$ and that, under this identification, the map induced on $\text{Pic}_K^0(E) = E(K)$ is equal to the connecting homomorphism in the Kummer sequence. We thus find ourselves in the ideal situation described at the beginning of this section.

2-descent on double covers of \mathbb{P}^1 . Let C be given by the equation⁷ $u_3^2 = g(u_1)$, where $g \in K[u_1]$ is separable of degree $2m$ with $m \geq 2$. The assumption that the degree is even means that the map to \mathbb{P}^1 does not ramify above ∞ ; applying a suitable change of coordinates one may always ensure that this is the case⁸. The assumption that $m \geq 2$ is made to exclude consideration of genus 0 curves. Let $A = K[u_1]/g(u_1)$ and denote the image of u_1 in A by θ . Then A is the étale algebra corresponding to the G_K -set of ramification points; θ is the map sending a ramification point to its u_1 -coordinate.

To do a 2-descent one typically uses the function $u_1 - \theta \in \kappa(C \otimes_K A)^\times$. The divisor of this function is $\text{div}(u_1 - \theta) = 2[(\theta, 0)] - ([\infty_+] + [\infty_-])$, where ∞_\pm denote the points on C lying above $\infty \in \mathbb{P}^1$. Note that $([\infty_+] + [\infty_-])$ is a sum of closed points of C . In the notation of the proposition, Ω is the set of ramification points, $\Psi = \{2\omega : \omega \in \Omega\}$ and their étale algebras are both isomorphic to A . The induced norm is given by squaring. So the proposition gives a homomorphism

$$\text{Pic}(C) \rightarrow \frac{A^\times}{K^\times A^{\times 2}}$$

If $\text{Pic}_K(C) = \text{Pic}(C)$, then as above this yields a homomorphism

$$J(K)/2J(K) \rightarrow A^\times/K^\times A^{\times 2},$$

where J is the Jacobian of C . This can be used to obtain information on the Mordell-Weil group. Since $C(K) \subset \text{Pic}_K^1(C)$, the homomorphism can also be used to study the rational points on C .

⁷The variables are labeled here so as to be compatible with the notation used in chapter II.

⁸If there happens to be some K -rational ramification point, we can arrange for it to lie above $\infty \in \mathbb{P}^1$. In this situation $\text{div}(u_1 - \theta) = 2[(0, \theta)] - 2[\infty_C]$ and the proposition yields a homomorphism to $A^\times/A^{\times 2}$

One can of course use other functions as well. For example the function $u_3 \in \kappa(C)^\times$ has divisor $\operatorname{div}(u_3) = \sum_{\omega \in \Omega} \omega - m(\infty_+ + \infty_-)$. Now the relevant algebras are A and K and the induced norm map is the usual norm $N_{A/K} : A \rightarrow K$. The proposition gives a homomorphism

$$\operatorname{Pic}(C) \rightarrow \frac{K^\times}{K^{\times m} N_{A/K}(A^\times)}.$$

The proposition applies as well to the pair $(u_1 - \theta, u_3)$ to give a map

$$\operatorname{Pic}(C) \rightarrow \frac{A^\times \times K^\times}{\iota(K^\times) \partial(A^\times)},$$

where $\iota : a \mapsto (a, a^m)$ and $\partial : \alpha \mapsto (\alpha^2, N_{A/K}(\alpha))$.

Unlike the situation for p -descent on elliptic curves, the map on $J(K)/2J(K)$ induced by $u_1 - \theta$ is not usually injective. In the second part of this thesis we will see how combining with the information from u_3 can be used to correct for this.

REMARK: The discussion here for hyperelliptic curves applies with minor changes to the more general case of cyclic covers of the projective line of the form $u_3^p = g(u_1)$ with $\deg(g)$ divisible by p . A detailed description of how the function u_3 can be used to similar effect in this context is the subject of a forthcoming paper of Stoll and Van Luijk [StVL].

Proof of Proposition 3.1. Any rational function $h \in \kappa(\bar{C})^\times$ defines a homomorphism from the group of divisors of C with support disjoint from the support of $\operatorname{div}(h)$ to the multiplicative group of \bar{K} by

$$d = \sum n_P P \mapsto h(d) = \prod h(P)^{n_P} \in \bar{K}^\times.$$

If K' is some extension of K and h is defined over K' , then this restricts to give a homomorphism from the group of K' -rational divisors with support disjoint from that of $\operatorname{div}(h)$ into K'^\times . If f is as in the proposition, then it is defined over $A(\Psi)$ which splits as a product of extensions of K , so the first statement in the proposition is clear.

For the second, define

$$\phi_f : \operatorname{Pic}(C) \rightarrow \frac{A(\Psi)^\times}{\iota(K^\times) \partial(A(\Omega)^\times)}$$

by setting the value of ϕ_f on $\Xi \in \operatorname{Pic}(C)$ equal to the class of $f(d)$, where $d \in \operatorname{Div}(C)$ is any K -rational divisor representing Ξ with support disjoint from Ω and $[f]_\infty$. If this is well-defined, then it is clearly the unique homomorphism with the stated property.

First we argue that such d exists. This follows from [La, page 166] where it is shown that any K -rational divisor class which is represented by a K -rational divisor contains a K -rational divisor avoiding a given finite set (see also [Sik3, footnote to page 4]). Next we use Weil reciprocity to show that the result does not depend on the choice for d .

Let $h \in \kappa(C)^\times$ be any rational function whose zeros and poles are disjoint from those of all of the f_ψ . We will show that $f(\operatorname{div}(h)) \in \iota(K^\times) \partial(A(\Omega)^\times)$, from which the proposition follows. For each $\psi \in \Psi$, the divisor of h is prime to

$$\operatorname{div}(f_\psi) = [f]_0(\psi) - [f]_\infty(\psi) = \psi - [f]_\infty(\psi).$$

So by Weil reciprocity,

$$f_\psi(\operatorname{div}(h)) = h(\operatorname{div}(f_\psi)) = \frac{h([f]_0(\psi))}{h([f]_\infty(\psi))}.$$

Interpreting this as a map we have

$$f(\operatorname{div}(h)) = \frac{h([f]_0)}{h([f]_\infty)} \in \operatorname{Map}_K(\Psi, \bar{K}^\times) = A(\Psi)^\times.$$

Define $\alpha \in \operatorname{Map}_K(\Omega, \bar{K}^\times) = A(\Omega)^\times$ by $\alpha : \Omega \ni \omega \mapsto h(\omega) \in \bar{K}^\times$. Now consider $\partial(\alpha) \in \operatorname{Map}_K(\Psi, \bar{K}^\times) = A(\Psi)^\times$. The value of $\partial(\alpha)$ at $\psi = \sum n_\omega \omega \in \Psi$ is

$$\partial(\alpha)_\psi = \prod \alpha(\omega)^{n_\omega} = \prod h(\omega)^{n_\omega} = h(\psi) = h([f]_0(\psi)).$$

This shows that $h([f]_0) = \partial(\alpha) \in \partial(A(\Omega)^\times)$.

It remains to show that $h([f]_\infty) \in \iota(K^\times)$. Recall that the value of $[f]_\infty$ on the orbit $\mathcal{O} \subset \Psi$ is the divisor $d_{\mathcal{O}} = \sum_{\mathcal{P}} m_{\mathcal{O}, \mathcal{P}} \mathcal{P}$ and that $m_{\mathcal{O}} = \gcd(m_{\mathcal{O}, \mathcal{P}})$. The \mathcal{P} are closed points on C . In particular, each is a K -rational divisor and we know how to evaluate h at \mathcal{P} to obtain an element in K^\times . Extending by linearity we have that the value taken by $h([f]_\infty)$ on any G_K -orbit $\mathcal{O} \subset \Psi$ is $\prod_{\mathcal{P}} h(\mathcal{P})^{m_{\mathcal{O}, \mathcal{P}}}$, which is a product of $m_{\mathcal{O}, \mathcal{P}}$ -th powers in K^\times . A product of $m_{\mathcal{O}, \mathcal{P}}$ -th powers is clearly a $\gcd(m_{\mathcal{O}, \mathcal{P}})$ -th power, so the value of $h([f]_\infty)$ on \mathcal{O} is in $K^{\times m_{\mathcal{O}}}$. It follows that $h([f]_\infty) \in \iota(K^\times)$. This completes the proof. \square

REMARK: Though we will make no use of it here, we point out that one can do slightly better in the last paragraph of the proof. For a closed point \mathcal{P} , use $K_{\mathcal{P}}$ to denote the residue field. Then in fact $h(\mathcal{P}) \in N_{K_{\mathcal{P}}/K}(K_{\mathcal{P}}^\times) \subset K^\times$. This means that one may be able to replace $\iota(K^\times)$ with a proper subgroup. The resulting homomorphism will then carry more information. For example, suppose C is the double cover of \mathbb{P}^1 defined by

$$u_3^2 = a_{2m} u_1^{2m} + \cdots + a_1 u_1 + a_0,$$

where $a_{2m} \in K^\times \setminus K^{\times 2}$ and set $L = K(\sqrt{a_{2m}})$. Then L is the residue field of the closed point at ∞ . The homomorphism induced by the typical $u_1 - \theta$ map used to do a 2-descent factors as:

$$\operatorname{Pic}(C) \longrightarrow \frac{A^\times}{N_{L/K}(L^\times)A^{\times 2}} \longrightarrow \frac{A^\times}{K^\times A^{\times 2}}.$$

It would be interesting to see if this can also be used to eliminate the ambiguity.

4. How to do a p -descent on an elliptic curve

Let E be an elliptic curve over a number field k and p a prime. In this section we provide the details of how one can compute the p -Selmer group of E using the ideas of the previous section. For odd p this was developed in [DSS, SchSt]. It is far more efficient than the naive algorithm delivered by the proof of theorem 2.2. The claim to efficiency stems from the fact that rather than working with the p -division field of E , one works with the algebra obtained by adjoining coordinates of a generic p -torsion point. Typically this is a field of degree $p^2 - 1$ while the p -division field is of degree $(p^2 - 1)(p^2 - p)$. Since the arithmetic computations required have exponential complexity in this degree this gives a significant improvement. With the current state of the art this gives a practical algorithm for performing 2-, 3- and perhaps even 5-descents on elliptic curves of reasonable size and over number fields of moderate degree and

discriminant.

We work again over K , a perfect field of characteristic not equal to p . Let $A = \text{Map}_K(E[p] \setminus \{0_E\}, \bar{K})$ be the étale K -algebra corresponding to the nontrivial p -torsion points. Let $f = (f_P) \in \kappa(C \otimes_K A)^\times$ be the function described in the example of the previous section. This induces a homomorphism $\Phi := \text{Pic}_K(E) \rightarrow A^\times/A^{\times p}$. Restricting to the degree 0 part we get a homomorphism $\Phi : \text{Pic}_K^0(E) = E(K) \rightarrow A^\times/A^{\times p}$. We now describe how this relates to the connecting homomorphism of the Kummer sequence.

Any map $\phi \in \bar{A} = \text{Map}(E[p] \setminus \{0_E\}, \bar{K})$ may be extended to a map defined on all of $E[p]$ by setting $\phi(0_E) = 1$. In this way the p -th roots of unity in \bar{A} may be identified with the G_K -set of maps from $E[p]$ to μ_p that take the value 1 at 0_E . Among such maps there are the homomorphisms, $\text{Hom}(E[p], \mu_p)$, which form a G_K -subset. Using the Weil pairing, $\text{Hom}(E[p], \mu_p)$ may be identified with $E[p]$. Taken together, these identifications give rise to an injection

$$0 \rightarrow E[p] \xrightarrow{w} \mu_p(\bar{A}).$$

By a generalization of Hilbert's Theorem 90, $H^1(K, \mu_p(\bar{A})) \simeq A^\times/A^{\times p}$, so there is an induced map $H^1(K, E[p]) \xrightarrow{w_*} A^\times/A^{\times p}$.

LEMMA 4.1. *The following diagram commutes.*

$$\begin{array}{ccc} E(K) & \xrightarrow{\Phi} & A^\times/A^{\times p} \\ & \searrow \delta & \nearrow w_* \\ & H^1(K, E[p]) & \end{array}$$

PROOF: This is essentially a cocycle computation. See [Sch2, Theorem 2.3] \square

As mentioned at the beginning of the previous section, for this to be useful one needs to know that w_* is injective and be able to describe its image. The map in question arises from the exact sequence

$$0 \rightarrow E[p] \xrightarrow{w} \mu_p(\bar{A}) \xrightarrow{q} Q \rightarrow 0,$$

where Q is used to denote the quotient of $\mu_p(\bar{A})$ by the image of $E[p]$. Taking Galois cohomology we have the exact sequence

$$\mu_p(A) \xrightarrow{q} H^0(K, Q) \rightarrow H^1(K, E[p]) \xrightarrow{w_*} A^\times/A^{\times p} \xrightarrow{q_*} H^1(K, Q).$$

So the kernel of w_* is the finite group $H^0(K, Q)/q(\mu_p(A))$ and the image of w_* is equal to the kernel of q_* . The strategy for describing both of these groups is to develop a more concrete description of the map $\mu_p(\bar{A}) \xrightarrow{q} Q$. This amounts to writing down a G_K -module morphism defined on $\text{Map}(E[p], \mu_p)$ with kernel equal to the subspace consisting of maps that are homomorphisms.

The case $p = 2$. Since we are in characteristic 2, a map $\phi : E[2] \rightarrow \mu_2$ is a homomorphism if and only if

$$\phi(P)\phi(Q)\phi(P+Q) = \phi(0_E) = 1, \text{ for some (any) basis } \{P, Q\} \text{ of } E[2].$$

For a map $\phi \in \mu_2(\bar{A})$, this condition simply means that $N(\phi) = 1$, where $N = N_{A/K} : \bar{A} \rightarrow \bar{K}$ is the norm. This gives a short exact sequence

$$0 \rightarrow E[2] \xrightarrow{w} \mu_2(\bar{A}) \xrightarrow{N} \mu_2 \rightarrow 0.$$

Hence, taking Galois cohomology and applying Hilbert's Theorem 90, we have an exact sequence⁹

$$1 \rightarrow \frac{\mu_2}{N(\mu_2(A))} \rightarrow H^1(K, E[2]) \xrightarrow{w_*} A^\times/A^{\times 2} \xrightarrow{N} K^\times/K^{\times 2}.$$

This allows us to determine the kernel of w_* . Since the action of G_K on μ_2 is trivial, $\mu_2(A)$ consists of those maps $(E[2] \setminus \{0_E\}) \rightarrow \mu_2$ which take constant values on the G_K -orbits. Since the number of non-trivial 2-torsion points is three, there must always be some orbit of odd order. It is then clear that the norm $N : \mu_2(A) \rightarrow \mu_2$ is surjective, and hence that w_* is injective.

From the exact sequence it is also clear that the image of w_* can be identified with the kernel of the norm on $A^\times/A^{\times 2}$. Taken together this shows that w_* gives an isomorphism

$$H^1(K, E[2]) \simeq \ker \left(N : \frac{A^\times}{A^{\times 2}} \rightarrow \frac{K^\times}{K^{\times 2}} \right).$$

The case $p \geq 3$. Assume now that p is an odd prime. The norm condition above no longer describes the homomorphisms completely. Clearly any homomorphism is homogeneous of degree 1, i.e. $\phi(nP) = n\phi(P)$ for all $n \in \mathbb{Z}$ and $P \in E[p]$. Recall that $E[p]$ is a 2-dimensional \mathbb{F}_p -vector space and that the action of G_K on $E[p]$ and all G_K -modules derived from it factors through $\mathrm{GL}_2(\mathbb{F}_p)$ (one says such modules have GL_2 -action). Being homogeneous of degree 1 means that the action of a central element $\alpha I \in \mathrm{GL}_2(\mathbb{F}_p)$ is multiplication by α . For a G_K -module M with GL_2 -action, use $M^{(1)}$ to denote the subgroup of elements that are homogeneous of degree 1. With this notation, we see that $w : E[p] \rightarrow \mu_p(\bar{A})$ actually gives a map $w : E[p] \rightarrow \mu_p(\bar{A})^{(1)}$. One can show that

$$H^1(K, \mu_p(\bar{A})^{(1)}) = \ker \left(g - \sigma_g : A^\times/A^{\times p} \rightarrow A^\times/A^{\times p} \right),$$

where g is a primitive root mod p and σ_g denotes the map induced on A by the action of g on $E[p]$ (see [SchSt, 5.2]). Thus the image of $w_* : H^1(K, E[p]) \rightarrow A^\times/A^{\times p}$ is contained in this kernel as well.

Now suppose $\phi : E[p] \rightarrow \mu_p$ is homogeneous of degree 1. In order that ϕ be a homomorphism it is necessary and sufficient that $\phi(P)\phi(Q) = \phi(P+Q)$ for all $P \in E[p] \setminus \{0_E\}$ and all $Q \in E[p] \setminus \{nP : n = 1, \dots, p\}$. Attempting to encode this in terms of a norm would lead to consideration of the algebra corresponding to the set of unordered triples of nontrivial p -torsion points that sum to 0_E but are not contained in a line through 0_E . This is an algebra of degree $\frac{1}{6} \# \mathrm{GL}_2(\mathbb{F}_p) = \frac{1}{6}(p^2 - 1)(p^2 - p)$. The following lemma characterizes homomorphisms in a way which allows us to get away with something much smaller.

LEMMA 4.2. *Let p be an odd prime and $\phi : \mathbb{F}_p^2 \rightarrow \mathbb{F}_p$ any map. For ϕ to be \mathbb{F}_p -linear the following two conditions are necessary and sufficient.*

- (1) $\forall a \in \mathbb{F}_p, x \in \mathbb{F}_p^2, \phi(ax) = a\phi(x)$ (i.e. ϕ is homogeneous of degree 1).
- (2) $\sum_{x \in \ell} \phi(x) = 0$, for any affine line $\ell \subset \mathbb{F}_p^2$ not passing through the origin.

PROOF: The necessity of these conditions is obvious. For the sufficiency see [SchSt, 5.7] \square

⁹One can check that the map induced by N after applying Theorem 90 is indeed the norm.

REMARK: When $p = 3$, the homogeneity is equivalent to requiring that the sum of the values on any line through the origin be 0. So together, conditions (1) and (2) amount to requiring that $\sum_{x \in \ell} \phi(x) = 0$ hold for all lines $\ell \subset E[3]$.

Since the action of G_K on $E[p]$ is linear, the set of all affine lines in $E[p]$ is stable under the action of Galois. Since 0_E is fixed, the subset \mathcal{L} of those lines not passing 0_E is also G_K -stable. Note that \mathcal{L} is derived from $E[p] \setminus \{0_E\}$ as a G_K -set in the sense of section 3. Let $B = \text{Map}_K(\mathcal{L}, \bar{K})$ be the étale algebra corresponding to \mathcal{L} . The induced norm map (coming from the structure of \mathcal{L} as a derived G_K -set),

$$N' : A \ni \phi \mapsto \left(\ell \mapsto \prod_{P \in \ell} \phi(P) \right) \in B,$$

encodes the second condition in 4.2. Namely, a map $\phi \in \mu_p(\bar{A})^{(1)}$ is a homomorphism if and only if $N'(\phi) = 1$. Thus one has an exact sequence

$$0 \rightarrow E[p] \rightarrow \mu_p(\bar{A})^{(1)} \xrightarrow{N'} \mu_p(\bar{B})^{(1)}.$$

Schaefer and Stoll have proven that this sequence remains exact when $H^1(K, -)$ is applied [SchSt, Proposition 5.8]. So there is an exact sequence

$$H^1(K, E[p]) \rightarrow H^1(K, \mu_p(\bar{A})^{(1)}) \xrightarrow{N'} H^1(K, \mu_p(\bar{B})^{(1)}).$$

Hence the image of $H^1(K, E[p])$ in $A^\times/A^{\times p}$ is equal to

$$\ker\left(g - \sigma_g : A^\times/A^{\times p} \rightarrow A^\times/A^{\times p}\right) \cap \ker\left(N' : A^\times/A^{\times p} \rightarrow B^\times/B^{\times p}\right).$$

REMARK: There are a total of $p(p+1)$ lines (in $E[p]$), with exactly $p+1$ passing through any given point, from which it follows that $\#\mathcal{L} = p^2 - 1$. So B is a product of fields of degree at most $p^2 - 1$ over K . To perform second p -descents we need a norm condition to cut out a space of affine maps. Since these maps no longer have a fixed point, we no longer have a reasonable notion of homogeneity. This forces us to use a larger algebra analogous to that mentioned preceding lemma 4.2 (see section II.2).

The proof that w_* is injective is somewhat more involved than in the case $p = 2$. The action of G_K on $E[p]$ corresponds to some conjugacy class of subgroups in $\text{GL}_2(\mathbb{F}_p)$. Essentially one classifies these and shows for each that the kernel of the induced map is trivial. The details can be found in [SchSt, DSS].

REMARK: To define $w_* : H^1(K, E[n]) \rightarrow A^\times/A^{\times n}$ there is no need to assume n is prime. One can show that, when n is any odd integer, w_* is injective provided the image of the representation $G_K \rightarrow \text{GL}_2(\mathbb{Z}/n\mathbb{Z})$ giving the action on $E[n]$ has large enough image [Cr1]. In the number field case, the map may fail to be injective somewhere locally even if it is globally injective. So this is not particularly well suited to computing the Selmer group. In general, there is an embedding into a certain quotient of $(A \otimes_K A)^\times$ [CFOSS-I, 3.2]. Generically this algebra will contain a number field of degree $O(n^4)$ over k , making computation of the unramified subgroup impractical. This embedding is useful, however, when one wants to find explicit models as coverings in projective space (see [CFOSS-I, -II]).

Computing the unramified subgroup. Now specialize to the case that $K = k$ is a number field. If we identify $H^1(k, E[p])$ with its image in $A^\times/A^{\times p}$ then for any set of primes S , we have

$$H^1(k, E[p]; S) = H^1(k, E[p]) \cap A(S, p).$$

In the case $p = 2$ this may be computed as the kernel of $N_{A/k} : A(S, 2) \rightarrow k(S, 2)$ which, after having computed $A(S, 2)$ and $k(S, 2)$ using theorem 2.3, can be reduced to linear algebra over \mathbb{F}_2 . For odd p one needs to compute the kernels of $g - \sigma_g$ and N' on $A(S, p)$. If one also computes $B(S, p)$ (actually it suffices to deal with $B(\emptyset, p)$ - see [SchSt, Section 7.1]), then the kernel of N' can be computed using linear algebra over \mathbb{F}_p . Alternatively one can compute $\ker(N')$ by checking directly which elements in the image are p -th powers in B . At least in the case $p = 3$, the map $g - \sigma_g$ coincides with the norm to a certain subalgebra. As such it is also amenable to computation via linear algebra. In any event, computation of $A(S, p)$ will dominate the running time. This leads to the following result, already alluded to at the beginning of the previous section.

THEOREM 4.3. *Let E be an elliptic curve over k , p a prime and S a finite set of primes (containing all infinite primes if $p = 2$). The cohomology group $H^1(k, E[p])$ embeds in $A^\times/A^{\times p}$ and there is an effective procedure for computing $H^1(k, E[p]; S)$ (as a subgroup of $A^\times/A^{\times p}$) which is efficient modulo computation of $A(S, p)$.*

From the unramified subgroup to the Selmer group. We know the p -Selmer group is contained in $H^1(k, E[p]; S)$ for an appropriate finite set of primes S . The next result gives a more precise statement. It also has the practical benefit of showing that we can usually take S to be a rather small set of primes.

THEOREM 4.4. *Let S be the set of primes of k containing all primes above p , all primes v such that the Tamagawa number of E at v is divisible by p , and all archimedean primes if $p = 2$. Then*

$$\text{Sel}^{(p)}(E/k) = \{ \xi \in H^1(k, E[p]; S) \mid \forall v \in S, \text{res}_v(\xi) \in \delta_v(E(k_v)) \}.$$

This is [SchSt, Proposition 3.2]. We will need a similar statement when we consider second p -descents. The ingredients of the proof will be the same; once one knows that the Selmer group is unramified outside S , the result is obtained from a counting argument outlined in the next two lemmas.

LEMMA 4.5. *Let v be a non-archimedean prime of k . Then*

$$\#E(k_v)/pE(k_v) = p^d \cdot \#E(k_v)[p],$$

where $d = [k_v : \mathbb{Q}_p]$ if v lies over p and $d = 0$ otherwise.

PROOF: Using the theory of formal groups, one shows that $E(k_v)$ contains a finite index subgroup isomorphic to the maximal ideal in the ring of integers of k_v [Mat]. This allows one to compute the size of the cokernel of multiplication by p . For details see [Sch1, Lemma 3.8] or [Sch2, Proposition 2.4]. \square

On the other hand, the size of the unramified subgroup of $H^1(k_v, E[p])$ can be computed using the inflation-restriction sequence.

LEMMA 4.6. *If v is a non-archimedean prime of k not lying over p , then the size of the unramified subgroup of $H^1(k_v, E[p])$ is equal to $\#E(k_v)[p]$.*

PROOF: This is shown in the proof of [Sch1, Lemma 3.1] \square

Since the connecting homomorphism induces an injective map on $E(k_v)/pE(k_v)$ we have the following.

COROLLARY 4.7. *For primes $v \notin S$, the unramified subgroup of $H^1(k_v, E[p])$ is equal to the image of the connecting homomorphism.*

Since we can compute $H^1(k, E[p]; S)$ using theorem 4.3, the problem of computing the p -Selmer group is reduced to checking, at the finitely many primes of S , which of the finitely many elements of $H^1(k, E[p]; S)$ restrict into the image of $E(k_v)$ under the connecting homomorphism $\delta_v : E(k_v) \rightarrow H^1(k_v, E[p])$. This is a finite problem. The explicit presentation of the connecting homomorphism in terms of the function $f = (f_P)$ makes it tractable.

Computing the local images. In practice, one makes use of the fact that A splits as a product of number fields. If $\{P_1, \dots, P_r\}$ are a set of representatives for the Galois orbits in $E[p] \setminus \{0_E\}$, then $A \simeq \prod_i k(P_i)$, where $k(P_i)$ denotes the extension of k obtained by adjoining the coordinates of P_i . The map $f = (f_P)$ is then determined by its values at P_i for $i = 1, \dots, r$. So it is enough to choose such a generating set and at each P_i a rational function f_{P_i} defined over $k(P_i)$, with a zero of order p at P_i and a pole of order p at 0_E .

The functions f_{P_i} also induce local maps which allow us to determine the image of δ_v as a subgroup of $A_v^\times/A_v^{\times p}$, where $A = A \otimes k_v$. At each prime v we have a commutative diagram of finite dimensional \mathbb{F}_p -vector spaces (we have identified $E(k)$ and $\text{Pic}_k^0(E)$).

$$\begin{array}{ccc} E(k)/pE(k) & \xrightarrow{f} & A(S, p) \\ \downarrow & & \text{res}_v \downarrow \\ E(k_v)/pE(k_v) & \xrightarrow{f_v} & A_v^\times/A_v^{\times p} \end{array}$$

The size of the lower left space is given by lemma 4.5. The horizontal maps are injective, hence this is also the size of the image of f_v . To compute it, it is enough to find the images of sufficiently many independent elements of $\text{Pic}_{k_v}^0(E)$ under f_v . The group in the lower right can be determined using Hensel's lemma and knowledge of the decomposition of v in the constituent fields of A . To determine the class of any $\alpha \in A_v^\times$ modulo p -th powers it is sufficient to know α up to some finite precision (see III.2.2). So in practice it is usually easier to determine independence by looking at the images in $A_v^\times/A_v^{\times p}$.

Identifying $H^1(k_v, E[p])$ with its image in $A_v^\times/A_v^{\times p}$, the image of f_v is the image of the local connecting homomorphism. Having computed it, its preimage in $A(S, p)$ can be found using linear algebra. Intersecting these preimages with $H^1(k, E[p]; S) \subset A(S, p)$, as v runs over S , yields the p -Selmer group. All of these local computations can be performed very quickly, whence the following result.

THEOREM 4.8. *Let E be an elliptic curve over k . There is an effective procedure for computing (a set of elements in A^\times representing) the p -Selmer group of E which is efficient modulo computation of $A(S, p)$.*

5. The interpretation as n -coverings

In the preceding sections we have dealt with the n -Selmer group as a purely algebraic object. In this section we discuss a geometric interpretation of the n -Selmer group elements as coverings of the elliptic curve. In addition to the conceptual advantages offered by a geometric perspective, the ability to represent n -Selmer group elements as geometric objects yields two important practical benefits. Firstly, it allows searching for points on the coverings. Provided one is able to produce nice models for the coverings, this should make finding points of large height easier. Secondly, it opens up the possibility of performing higher descents. This, the primary focus of the thesis, is taken up in the next two chapters.

The twisting principle. The connection between the Galois cohomology groups of the preceding sections and the geometric objects we seek is given by the so-called twisting principle. If \mathcal{A} is some algebro-geometric object defined over K , then a twist of \mathcal{A} is an object \mathcal{A}' , also defined over K , which is isomorphic to \mathcal{A} over some algebraic extension of K . Two twists are said to be (K -)isomorphic if they are already isomorphic over K . The twisting principle states that the set of isomorphism classes of twists of \mathcal{A}/K is classified by the pointed set $H^1(K, \text{Aut}(\mathcal{A}))$ (The distinguished point corresponding to the isomorphism class of \mathcal{A}). The validity of this ‘principle’ depends, of course, on the objects considered.

If \mathcal{A}' is a twist of \mathcal{A} , then there is an isomorphism $\phi : \mathcal{A}' \rightarrow \mathcal{A}$ defined over some extension of K . One associates to this the map

$$G_k \ni \sigma \mapsto \phi^\sigma \phi^{-1} \in \text{Aut}(\mathcal{A}).$$

One can check that this is a cocycle and that its class in $H^1(K, \text{Aut}(\mathcal{A}))$ does not depend on the choice for ϕ . This gives an injective map from the set of K -isomorphism classes of twists into $H^1(K, \text{Aut}(\mathcal{A}))$. If \mathcal{A} is a quasi-projective variety then this map is also surjective, and hence an isomorphism [Ser2, Ch. V, Cor. 2 to Prop. 12]. Thus given a cocycle $\xi \in H^1(K, \text{Aut}(\mathcal{A}))$, one may form the twist of \mathcal{A} by ξ .

Following [CFOSS-I, Section 1] we claim that the map will also be surjective in all of our applications. Our objects will be quasi-projective K -varieties equipped with additional structure. To form the twist of such an object, one first forms the twist of the underlying quasi-projective variety. The additional structure may then be transferred using the resulting isomorphism.

Torsors. Let \mathcal{G} be a commutative algebraic group defined over K . A K -torsor (\mathcal{T}, μ) under \mathcal{G} is a twist \mathcal{T} of \mathcal{G} together with a simply transitive algebraic group action

$$\mu : \mathcal{G} \times \mathcal{T} \rightarrow \mathcal{T}$$

of \mathcal{G} on \mathcal{T} defined over K . An isomorphism of K -torsors \mathcal{T} and \mathcal{T}' under \mathcal{G} is an isomorphism $\mathcal{T} \simeq \mathcal{T}'$ which is compatible with the action of \mathcal{G} . The action of \mathcal{G} on itself by translation gives \mathcal{G} the structure of a K -torsor under \mathcal{G} and any K -torsor under \mathcal{G} which possesses a K -rational point is K -isomorphic to this trivial torsor. The automorphism group of this trivial torsor is isomorphic to \mathcal{G} . So by the twisting principle, $H^1(K, \mathcal{G})$ classifies the set of K -isomorphism classes of K -torsors under \mathcal{G} .

In this way, if E is an elliptic curve over K , then $H^1(K, E)$ classifies K -torsors under E and any class in $H^1(K, E)$ can be represented by some smooth, projective and

absolutely irreducible genus one curve T/K with a simply transitive, Galois equivariant action of E . This will be a trivial torsor precisely when $T(K) \neq \emptyset$. One can express the group law in $H^1(K, E)$ geometrically in terms of these genus one curves [Sil, Exer. 10.2]. With this interpretation, the group $H^1(K, E)$ is often referred to as the Weil-Châtelet group of E/K . Conversely, any smooth, projective and absolutely irreducible genus one curve T defined over K can be given the structure of a K -torsor under its Jacobian E . Up to isomorphism, the only possible ambiguity in the choice comes from the automorphisms of E as an elliptic curve. When the j -invariant of E is neither 0 nor 1728, the only automorphisms are $\{\pm 1\}$. In any case, T can be endowed with at most finitely many non-isomorphic structures of torsor under E . We will usually abuse notation by writing T for a K -torsor under E when in fact the action of E on T is part of the data.

Similarly $H^1(K, E[n])$ classifies isomorphism classes of K -torsors under $E[n]$. Such an object is a finite G_K -set together with a compatible action of $E[n]$ which is simply transitive. This gives a useful interpretation, but there are also others.

n -coverings. Let C and D be smooth, projective and absolutely irreducible curves defined over K and $\pi : D \rightarrow C$ a finite étale morphism defined over K . We say that (D, π) is a Galois covering if the group of \bar{K} -automorphisms of D considered as a scheme over C acts simply transitively on each fiber of π . In this situation we refer to the group of \bar{K} -automorphisms of the covering as the Galois group. There is a natural action of G_K on these automorphisms. If the Galois group is abelian, then it is a G_K -module and the notions of Galois covering of C with group M and of irreducible C -torsor under M are synonymous. The term M -covering will also be used to mean a Galois covering with group M .

If M is any finite G_K -module and $D \rightarrow C$ is an M -covering, then it follows from the twisting principle that the set of isomorphism classes of Galois coverings of C with group M is parameterized by $H^1(K, M)$. It is important to note that one must assume the existence of at least one such covering. Even then, there will in general be no canonical choice for a trivial covering. For this reason we interpret the set of isomorphism classes of M -coverings of C as a principal homogeneous space for $H^1(K, M)$ rather than as a group. The action is given, of course, by twisting. In what follows, we will primarily be interested in the case where $M = E[n]$.

DEFINITION 5.1. *Let C be a smooth, projective and absolutely irreducible curve defined over K with Jacobian E and $n \geq 2$ prime to the characteristic of K . An n -covering of C over K is a Galois covering of C defined over K with Galois group isomorphic to $E[n]$. We denote the set of all K -isomorphism classes of n -coverings of C defined over K by $\text{Cov}^{(n)}(C/K)$.*

Geometrically, every n -covering of C is obtained by pulling-back the multiplication by n map via a suitable embedding of C into the Jacobian. When $C = E$ is an elliptic curve, every n -covering of E can be viewed as a twist of the multiplication by n map on E . This gives a canonical choice for the trivial n -covering of E . So $\text{Cov}^{(n)}(E/K)$ is canonically isomorphic to $H^1(K, E[n])$, and as such is considered to be a group. In general (provided it is nonempty), we consider $\text{Cov}^{(n)}(C/K)$ as principal homogeneous space for $H^1(K, E[n])$ with the action being given by twisting.

Let us quickly reinterpret the Kummer sequence

$$0 \rightarrow E(K)/nE(K) \xrightarrow{\delta} H^1(K, E[n]) \rightarrow H^1(K, E)[n] \rightarrow 0$$

using the language of torsors and n -coverings. It is evident from the definition of the connecting homomorphism that a point $P \in E(K)$ is mapped to the class of the n -covering given by

$$\pi_P : E \ni Q \mapsto (nQ + P) \in E.$$

Note that the image of π_P on the Mordell-Weil group is the coset $\pi_P(E(K)) = P + nE(K)$. If (D, π) is an n -covering of E , then by definition (D, π) is a twist of (E, n) , so there is some isomorphism $\psi : D \rightarrow E$ defined over \bar{K} such that $\pi = n \circ \psi$. This is used to give D the structure of torsor under E via the rule

$$\mu(Q, x) = \psi^{-1}(\psi(x) + Q),$$

the addition taking place on E . The map $H^1(K, E[n]) \rightarrow H^1(K, E)[n]$ in the sequence is the forgetful map sending a covering (D, π) to the class of D (with the action of E via μ above) in $H^1(k, E)[n]$. A K -torsor under E is trivial precisely when it has a K -rational point, so the image of $E(K)/nE(K)$ in $H^1(K, E[n])$ consists precisely of those (isomorphism classes of) n -coverings which have K -rational points.

In particular, this shows that the n -coverings of E partition its K -rational points:

$$E(K) = \coprod_{(D, \pi) \in \text{Cov}^{(n)}(E/K)} \pi(D(K)).$$

The nonempty sets in this union are the cosets of $nE(K) \subset E(K)$. When $K = k$ is a number field, these cosets are finite in number. More generally we have the following fundamental theorem, which goes back to Chevalley and Weil [ChWe].

THEOREM 5.2. *Let C be a smooth, projective and absolutely irreducible curve defined over K and M a finite G_K -module. Suppose that over \bar{K} there exists an M -covering of C . Then every K -rational point $P \in C(K)$ is the image of a K -rational point on some M -covering of C defined over K , which is unique up to K -isomorphism. Moreover, if K is a number field or the completion of a number field, then only finitely many isomorphism classes of M -coverings are represented by curves with K -rational points.*

REMARK: The need to assume that a ‘geometric M -covering’ exists is forced by topology. In order that M -coverings exist, M must be a finite quotient of the (étale) fundamental group of C .

Composite coverings. Let $m, n \geq 2$ be integers not divisible by the characteristic of K , E an elliptic curve over K and (C, ρ) an n -covering of E . We would like to relate the sets $\text{Cov}^{(m)}(C/K)$ and $\text{Cov}^{(mn)}(E/K) = H^1(K, E[mn])$.

There is an exact sequence,

$$0 \rightarrow E[m] \xrightarrow{i} E[mn] \xrightarrow{m} E[n] \rightarrow 0.$$

From this we deduce an exact sequence

(5.1)

$$0 \rightarrow \frac{E(K)[n]}{mE(K)[mn]} \rightarrow H^1(K, E[m]) \xrightarrow{i_*} H^1(K, E[mn]) \xrightarrow{m_*} H^1(K, E[n]) \xrightarrow{\alpha} H^2(K, E[m]),$$

where α is a connecting homomorphism.

We call any $(D, \pi) \in H^1(K, E[mn])$ such that $m_*(D, \pi) = (C, \rho)$ a *lift* of (C, ρ) to an mn -covering. There is a canonical map $\text{Cov}^{(m)}(C/K) \rightarrow H^1(K, E[mn])$, given by composing the covering maps. If (D', π') is an m -covering of C , its image under this map is the class of $(D', \rho \circ \pi')$ which is an mn -covering of E . The image of this map is

the set of all lifts of (C, ρ) to an mn -covering. The fibers of this map are parameterized by the finite group $\frac{E(K)[n]}{mE(K)[mn]}$. In this way one reduces the study of mn -coverings of E to the study of m -coverings of the n -coverings of E .

One can further reduce the problem to the study of n -coverings where n is a prime power.

LEMMA 5.3. *Suppose m and n are relatively prime and not divisible by the characteristic of K . Then*

$$H^1(K, E[mn]) \simeq H^1(K, E[m]) \times H^1(K, E[n])$$

PROOF: Consider the exact sequence (5.1). Since $H^1(K, E[n])$ is of exponent n and m is prime to n , the cokernel of m_* vanishes. The same argument applies to the kernel of i_* . So (5.1) reduces to the required split exact sequence.

Alternatively one can use that fact that $E[mn]$ and $E[m] \times E[n]$ are isomorphic G_K -modules when m and n are relatively prime. \square

In some situations it may be possible that an n -covering of E have no lift to an mn -covering. The connecting homomorphism α in (5.1) gives the obstruction to finding a lift. If C is an everywhere locally solvable n -covering of E over a number field k , then this obstruction vanishes everywhere locally (since the existence of points implies the existence of a lift). Tate has shown [Ca2, Section 5] that when m is prime, the group $H^2(k, E[m])$ satisfies the Hasse principle. This means any class in $H^2(k, E[m])$ that is everywhere locally trivial is trivial. Whence the following result.

THEOREM 5.4 (Tate). *Suppose E is an elliptic curve over a number field k and (C, ρ) is an everywhere locally solvable n -covering of E defined over k . Then, for any prime number p , the set $\text{Cov}^{(p)}(C/k)$ is non-empty.*

REMARK: This result implies that every element of $\text{III}(E/k)$ is divisible by p in $H^1(k, E)$ for every prime p .

The Selmer set. Let C be a smooth, projective and absolutely irreducible curve defined over a number field k . In order that $C(k) \neq \emptyset$, C must have a point over every completion. For curves of genus 0 it is well known that this is already sufficient. It is equally well known that this does not hold in general. Theorem 5.2 provides another obstruction to the existence of k -points: In order that $C(k) \neq \emptyset$ it is necessary (and sufficient) that for any suitable finite G_k -module M there exist an M -covering of C defined over k with a k -rational point. Such a covering must necessarily be everywhere locally solvable. With this in mind, we make the following definition.

DEFINITION 5.5. *Let C be a smooth, projective and absolutely irreducible curve defined over a number field k . We define the n -Selmer set of C/k to be the set of all isomorphism classes of n -coverings of C over k that are everywhere locally solvable:*

$$\text{Sel}^{(n)}(C/k) = \{(D, \pi) \in \text{Cov}^{(n)}(C/k) \mid \text{for all primes } v, D(k_v) \neq \emptyset\}.$$

It is immediate that $(\text{Sel}^{(n)}(C/k) = \emptyset) \Rightarrow (C(k) = \emptyset)$.

When $C = E$ is an elliptic curve, $\text{Sel}^{(n)}(E/k)$ is a subgroup of $\text{Cov}^{(n)}(E/k) = H^1(k, E[n])$. We already defined $\text{Sel}^{(n)}(E/k)$ in section 2. Of course the two definitions

are equivalent: The condition that a class $\xi \in H^1(k, E[n])$ restrict into the local image of the connecting homomorphism is equivalent to requiring that every n -covering representing ξ have a local point.

Similarly, the Tate-Shafarevich group $\text{III}(E/k)$ may be interpreted as the group of k -isomorphism classes of k -torsors under E that have points everywhere locally. As such, nontrivial elements are represented by smooth genus one curves over k which give counter-examples to the Hasse principle.

We can interpret the method for p -descent described in the previous section from this geometric perspective as well. Let S be the set of bad primes for E given in theorem 4.4. Then $H^1(k, E[p]; S)$ is the finite set of isomorphism classes p -coverings of E which have a k_v point at all primes $v \notin S$. To cut out the Selmer set one needs to determine which of these are also locally solvable at the primes in S . This was achieved algebraically by using the explicit description of the connecting homomorphism in terms of the function $f = (f_P)$; In principle it could also be accomplished by determining local solvability of the coverings.

Second descents. Performing a n -descent on E one obtains the finite group $\text{Sel}^{(n)}(E/k)$. Determining the subgroup $E(k)/nE(k) \subset \text{Sel}^{(n)}(E/k)$ is equivalent to deciding which of these n -coverings have rational points. If $(C, \pi) \in \text{Sel}^{(n)}(E/k)$, then there is no local obstruction to the existence of a k -rational point on C . But this is not the only obstruction; C may fail to have a rational point because $\text{Sel}^{(n)}(C/k) = \emptyset$. A second n -descent determines the subgroup of $\text{Sel}^{(n)}(E/k)$ consisting of coverings for which this more refined obstruction is also trivial.

DEFINITION 5.6. *To do a second n -descent on E means to compute $\text{Sel}^{(n)}(C/k)$ for some (multiple, all) covering(s) $(C, \pi) \in \text{Sel}^{(n)}(E/k)$.*

It is clear from the definitions that the map

$$H^1(k, E[n^2]) \xrightarrow{n^*} H^1(k, E[n]).$$

sends the n^2 -Selmer group to the n -Selmer group. So one possibility for computing the n^2 -Selmer group is to compute the fiber above each element in $\text{Sel}^{(n)}(E/k)$. Up to the equivalence furnished by the kernel in the exact sequence (5.1), $\ker(i_*) = \frac{E(k)[n]}{nE(k)[n^2]}$, this is the same as computing $\text{Sel}^{(n)}(C/k)$ for each $(C, \rho) \in \text{Sel}^{(n)}(E/k)$. Since $\frac{E(k)[n]}{nE(k)[n^2]}$ is finite and (easily) computable, the information one obtains this way is just as good. Thus one reduces the problem of performing n^2 -descents on elliptic curves to that of performing second n -descents.

REMARK: In the language of [Ca1], the n -coverings $(C, \rho) \in \text{Sel}^{(n)}(E/k)$ for which $\text{Sel}^{(n)}(C/k) \neq \emptyset$ are the coverings which ‘survive the second descent’. The classes of such C in $\text{III}(E/k)$ form the kernel of the Cassels-Tate pairing on $\text{III}(E/k)[n]$. For the purposes of bounding the rank of $E(k)$, this kernel and the second descent yield the same information. However, performing a second n -descent gives more information. Namely, one should be able to construct explicit models representing the coverings.

The following lemma shows that for the purpose of bounding the rank of the Mordell-Weil group or the order of the Shafarevich-Tate group one should not need to perform descents on every element in the Selmer group.

LEMMA 5.7. *Let E be an elliptic curve over a number field k , p a prime and suppose that $\text{III}(E/k)[p]$ has order $\leq p^2$. Then the following are equivalent.*

- (1) $\text{III}(E/k)[p^\infty] = \text{III}(E/k)[p] \simeq \mathbb{Z}/p\mathbb{Z} \times \mathbb{Z}/p\mathbb{Z}$.
- (2) *There exists $(C, \pi) \in \text{Sel}^{(p)}(E/k)$ such that $\text{Sel}^{(p)}(C/k) = \emptyset$.*

PROOF: Let $N = \# \left(\frac{\text{III}(E/k)[p]}{p \text{III}(E/k)[p^2]} \right)$ be the number of classes in the quotient of $\text{III}(E/k)[p]$ by the subgroup of elements divisible by p in $\text{III}(E/k)$. The Cassels-Tate pairing induces a nondegenerate alternating bilinear form on this quotient. It follows that N is a square. On the other hand, (1) is equivalent to requiring that $N = p^2$, while (2) is equivalent to having $N > 1$. The result follows since p is prime. \square

REMARK: Note also that when (1) holds there will be $(\#E(k)/pE(k)) \cdot (p^2 - 1)$ elements of the Selmer group satisfying (2), so such an example should not be difficult to find. One can obviously employ similar tricks when the rank of $\text{III}(E/k)[p]$ is larger or when p is replaced by a prime power.

6. Projective models

We now want to discuss the possibility of representing these abstractly defined geometric objects concretely as curves in projective space. We work over a perfect field K and fix integers $m, n \geq 2$ not divisible by the characteristic.

Period, index and genus one normal curves. Let T be a smooth, projective and absolutely irreducible genus one curve defined over K with Jacobian E . The period, $\text{per}(T)$, of T is the least positive degree of a K -rational divisor class on T . Equivalently, T can be endowed with the structure of a torsor under E and $\text{per}(T)$ is the order of this class in the group $H^1(K, E)$ (the order does not depend on the choice of structure). One defines the index, $\text{ind}(T)$, of T to be the least positive degree of a K -rational divisor on T . Clearly $\text{per}(T) \leq \text{ind}(T)$. We will say that T has a *period-index obstruction* if the two are not equal. Note also that $T(K) \neq \emptyset$ if and only if $\text{per}(T) = \text{ind}(T) = 1$.

For any K -rational divisor of degree $n \geq 2$ on T , the associated complete linear system gives rise to a morphism from T to \mathbb{P}^{n-1} defined over K . For $n = 2$ this results in a double cover of \mathbb{P}^1 ramified in four points, which gives an (affine) model for T of the form $y^2 = f(x)$ where $f \in K[x]$. One can arrange to have $\deg(f) = 4$ by changing coordinates so that the covering is not ramified above $\infty \in \mathbb{P}^1$. Such an object (or the normalized model in the $(1, 2, 1)$ -weighted projective plane) is called a genus one normal curve of degree 2. For $n \geq 3$, the complete linear system yields an embedding $T \hookrightarrow \mathbb{P}^{n-1}$. The image is called a genus one normal curve of degree n . For $n = 3$, the image of T is a plane cubic curve. For larger n , the homogeneous ideal of the image is generated by a K -vector space of quadrics of dimension $n(n-3)/2$. For the sake of completeness, we will say that a genus one normal curve of degree 1 is a curve given by a Weierstrass equation $y^2 = f(x)$ with f separable and of degree 3. One easily sees that T admits a model as a genus one normal curve of degree n if and only if $\text{ind}(T)$ divides n .

More generally, the complete linear system associated to any K -rational divisor class on T of degree $n \geq 2$ gives rise to a K -morphism $T \rightarrow S$ from T to a Brauer-Severi variety S of dimension $n-1$ (see [Ser1, p. 160], [CFOSS-I, 1.20], [Cl, Section 3]). Conversely, if $T \rightarrow S$ is such a morphism, then over \bar{K} we have that $S \simeq \mathbb{P}^{n-1}$. Assuming T is not contained in any hyperplane, the pull-back of the hyperplane class

is a K -rational divisor class on T of degree n . Such a morphism exists if and only if $\text{per}(T)$ divides n .

This leads to the notions of torsor divisor class pairs and Brauer-Severi diagrams. The data for a torsor divisor class pair consists of a K -torsor T under E and a K -rational divisor class of degree n on T . The corresponding morphism $T \rightarrow S$ is called a Brauer-Severi diagram. In [CFOSS-I, Section 1] it is shown that, up to appropriate notions of isomorphism, torsor divisor class pairs and Brauer-Severi diagrams are both parameterized by the group $H^1(K, E[n])$. Recall that this group also parameterizes n -coverings of E . If (C, ρ) is an n -covering, then there exists an isomorphism $\psi : C \rightarrow E$ defined over \bar{K} such that $\rho = n \circ \psi$. This gives C the structure of a K -torsor under E . The pull-back of $n[0_E]$ by ψ defines a K -rational divisor class on C . One can show that this gives a torsor divisor class pair, whose class in $H^1(K, E[n])$ is the same as that of (C, ρ) . Thus the Brauer-Severi diagram corresponding to (C, ρ) is the map $C \rightarrow S$ given by the complete linear system associated to the divisor $\psi^*n[0_E]$. This results in a model for C as a genus one normal curve of degree n in \mathbb{P}^{n-1} if and only if $\psi^*n[0_E]$ is linearly equivalent to some K -rational divisor.

The obstruction map. Recall that $\text{Pic}(T)$ is the quotient of the group of K -rational divisors on T by the group of K -rational principal divisors, while $\text{Pic}_K(T) = \text{Pic}(\bar{T})^{G_K}$ is the group of K -rational divisor classes. It follows from Hilbert's Theorem 90 that the obvious map $\text{Pic}(T) \rightarrow \text{Pic}_K(T)$ is injective. In general, however, it is not surjective. To measure this failure one is naturally led to use Galois cohomology. The Picard group is defined by the exact sequence

$$1 \rightarrow \bar{K}^\times \rightarrow \kappa(\bar{T})^\times \rightarrow \text{Div}(\bar{T}) \rightarrow \text{Pic}(\bar{T}) \rightarrow 0.$$

Taking Galois invariants, this sequence may no longer be exact. One can deduce an exact sequence

$$0 \rightarrow \text{Pic}(T) \rightarrow \text{Pic}_K(T) \xrightarrow{\delta_T} \text{Br}(K),$$

where $\text{Br}(K)$ denotes the Brauer group of K . The map δ_T gives the obstruction to a K -rational divisor class being defined by a K -rational divisor.

Following [CFOSS-I] we define the obstruction map

$$\text{Ob}_n : H^1(K, E[n]) \rightarrow \text{Br}(K)$$

by $\text{Ob}_n(\xi) = \delta_T(\Xi)$, where (T, Ξ) is any torsor divisor class pair representing the class $\xi \in H^1(K, E[n])$. From this definition we obtain the fundamental property of the obstruction map that the Brauer-Severi diagram corresponding to an n -covering (C, ρ) gives a model for C as a genus one normal curve of degree n in \mathbb{P}^{n-1} if and only if $\text{Ob}_n((C, \rho)) = 0$. Conversely, any genus one normal curve $T \rightarrow \mathbb{P}^{n-1}$ of degree n , together with a structure of torsor under E determines a unique isomorphism class of n -coverings of E with trivial obstruction. Recall also that the torsor structure is unique up to automorphisms of E (as an algebraic group).

O'Neil has shown [O'N] that the obstruction map is quadratic. This means that, for any integer a , $\text{Ob}_n(a\xi) = a^2 \text{Ob}_n(\xi)$ and that the pairing

$$(\xi, \xi') \mapsto \text{Ob}_n(\xi + \xi') - \text{Ob}_n(\xi) - \text{Ob}_n(\xi')$$

is bilinear. The pairing is in fact the cup product associated to the Weil pairing on $E[n]$, i.e. the composition

$$\cup_n : H^1(K, E[n]) \times H^1(K, E[n]) \xrightarrow{\cup} H^2(K, E[n] \otimes E[n]) \xrightarrow{e_n} H^2(K, \mu_n) \simeq \text{Br}(K)[n].$$

Extension to $\text{Cov}^{(n)}(C/K)$. We now allow $n \geq 1$ with the convention that a 1-covering is the identity map. Let (C, ρ) be an n -covering of E defined over K . We would like to extend the obstruction map to the set of isomorphism classes of m -coverings of C . To do this, we simply compose with the canonical map $\text{Cov}^{(m)}(C/K) \rightarrow \text{H}^1(K, E[mn])$ that is given by composing the covering maps. We denote the resulting map,

$$\text{Cov}^{(m)}(C/K) \rightarrow \text{H}^1(K, E[mn]) \xrightarrow{\text{Ob}_{mn}} \text{Br}(K),$$

also by Ob_{mn} . The reader is cautioned that this map and, consequently, the set in the following definition depend on the structure of C as an n -covering of E . This slight abuse of notation should cause no confusion.

DEFINITION 6.1. *We say that an m -covering $\pi : D \rightarrow C$ has trivial obstruction if its image under Ob_{mn} is trivial. We use*

$$\text{Cov}_0^{(m)}(C/K) := \{(D, \pi) \in \text{Cov}^{(m)}(C/K) : \text{Ob}_{mn}((D, \pi)) = 0\}$$

to denote the set of isomorphism classes of m -coverings of C with trivial obstruction.

The obstruction maps of levels m , n and mn are related in the following lemma. From this one sees that $\text{Cov}_0^{(m)}(C/K) = \emptyset$ if $\text{Ob}_n((C, \rho)) \neq 0$, and, in the particular case $m = n$, that $\text{Ob}_{n^2} \circ i_* = 0$.

LEMMA 6.2. *The following diagram commutes.*

$$\begin{array}{ccccc} \text{H}^1(K, E[m]) & \xrightarrow{i_*} & \text{H}^1(K, E[mn]) & \xrightarrow{m_*} & \text{H}^1(K, E[n]) \\ \downarrow \text{Ob}_m & & \downarrow \text{Ob}_{mn} & & \downarrow \text{Ob}_n \\ \text{Br}(K)[m] & \xrightarrow{n} & \text{Br}(K)[mn] & \xrightarrow{m} & \text{Br}(K)[n] \end{array}$$

PROOF: One can prove this using the compatibility of the Weil pairings of levels mn and n and the fact that the bilinear form associated to the obstruction map is the Weil pairing cup product. For details see [CS, Proposition 6]. \square

REMARK: The obstruction map for n -coverings is closely related to the period-index obstruction for the underlying curves. If T is a K -torsor under E of period dividing mn , then T has index dividing mn if and only if there is a map $\pi : T \rightarrow E$ making T into an mn -covering of E with trivial obstruction. In particular, if $\text{per}(C) = \text{ind}(C) = n$, then $\text{Cov}_0^{(m)}(C/K) \neq \emptyset$ implies that there is some $T \in \text{H}^1(K, E)$ with trivial period-index obstruction such that $mT = C$. However, if (D, π) is an m -covering of C and D has no period-index obstruction it is not in general true that (D, π) has trivial obstruction (even for $C = E$).

If D is a smooth, projective and absolutely irreducible curve over K and $D(K) \neq \emptyset$, then $\text{Pic}(D) = \text{Pic}_K(D)$. If K is a number field and D is everywhere locally solvable, then this is the case everywhere locally. From the exact sequence $\text{Pic}(D) \rightarrow \text{Pic}_K(D) \rightarrow \text{Br}(D)$ and the local-global principle for the Brauer group we see that the existence of points everywhere locally implies that $\text{Pic}(D) = \text{Pic}_K(D)$. This is a result of Cassels [Ca2, Theorem 1.2]. It follows that an m -covering $\pi : D \rightarrow C$ has trivial obstruction if $D(K) \neq \emptyset$ or if K is a number field and D is everywhere locally solvable. In particular, the elements of the m -Selmer set (resp. m -Selmer group if $C = E$) have trivial obstruction.

REMARK: Recall the exact sequence $E(K) \xrightarrow{\delta} H^1(K, E[n]) \rightarrow H^1(K, E)[n] \rightarrow 0$. Since the image of δ consists of those classes of n -coverings which have a K -rational point, it is tempting to think that Ob_n should factor through $H^1(K, E)[n]$. However, the obstruction map is not a homomorphism, so this need not be the case (cf. the previous remark).

The following lemma gives an alternative characterization of the obstruction map when $m = n$ which will be fundamental to our definition of the descent map in the next chapter.

LEMMA 6.3. *Let $\rho : C \rightarrow E$ be an n -covering defined over K , set $X = \rho^{-1}(0_E)$ and let $(D, \pi) \in \text{Cov}^{(n)}(C/K)$. Then (D, π) has trivial obstruction if and only if there exists a model for D as a genus one normal curve of degree n^2 in \mathbb{P}^{n^2-1} defined over K with the property that the pull-back of any $x \in X$ by π is a hyperplane section.*

PROOF: Fix isomorphisms $\psi_D : D \rightarrow E$ and $\psi_C : C \rightarrow E$ (defined over \bar{K}) such that the diagram

$$\begin{array}{ccccc} D & \xrightarrow{\pi} & C & \xrightarrow{\rho} & E \\ \psi_D \downarrow & & \psi_C \downarrow & & \parallel \\ E & \xrightarrow{n} & E & \xrightarrow{n} & E \end{array}$$

commutes. By definition (D, π) has trivial obstruction if and only if $\psi_D^*(n^2[0_E])$ is linearly equivalent to some K -rational divisor. On the other hand, D admits a model as in the statement of the lemma if and only if $\pi^*[x]$ is linearly equivalent to some K -rational divisor, for each $x \in X$. It thus suffices to show, for all $x \in X$, that $\psi_D^*(n^2[0_E])$ and $\pi^*[x]$ are linearly equivalent. For this we may work geometrically. The problem is then equivalent to showing that for any n -torsion point $P \in E[n]$, the pull-back of P under the multiplication by n isogeny is linearly equivalent to $n^2[0_E]$. This follows from the well-known fact that two divisors on an elliptic curve are linearly equivalent if and only if they have the same degree and the same sum. Indeed, the divisors in question both have degree n^2 and sum to 0_E in the group $E(\bar{K})$. \square

Computing models explicitly. Let E be an elliptic curve over K . We have the rather abstractly defined subset $\text{Cov}_0^{(n)}(E/K) \subset H^1(K, E[n])$ of isomorphism classes of n -coverings that admit a model as a genus one normal curve of degree n . On the other hand, when p is prime, we have a more concrete realization of $H^1(K, E[p])$ as a subgroup of $A^\times/A^{\times p}$, where A is the étale K -algebra associated to the non-trivial p -torsion points of E . Given $a \in A^\times$ representing a class in $\text{Cov}_0^{(p)}(E/K)$, one would like to be able to explicitly compute a set of defining equations.

The main result of [CFOSS-I, -II] is a method for doing just that. In fact they allow n to be an arbitrary integer greater than 2, the situation for 2-descent having been well-known for some time (e.g. [Ca4, Section 15]). In section 4 we remarked that for arbitrary n , $H^1(K, E[n])$ can be embedded in a quotient of the étale K -algebra $R^\times := \text{Map}_K(E[n] \times E[n], \bar{K}^\times)$. It is from this algebraic presentation that they compute models explicitly. When n is prime it is also shown how to obtain representatives in R^\times from the subgroup of $A^\times/A^{\times n}$ computed by the method described in section 4. So taken together these give a method for performing explicit p -descents.

For their method to work, they require access to a ‘Black Box’ which, given structure constants for a K -algebra known to be isomorphic to $\text{Mat}_n(K)$, computes an isomorphism explicitly. This allows them to ‘trivialize’ the obstruction algebra associated to a class in $\text{Cov}_0^{(n)}(E/K)$. For $n = 2$ this amounts to finding a point on a conic. For $n = 3$ and $K = \mathbb{Q}$ there is a practical method for performing the role of the ‘Black Box’ which is part of the 3-descent implementation in MAGMA. The details of this are to be discussed in [CFOSS-III].

We summarize their result in the following theorem. While this becomes the starting point for our second p -descents, the details of the method itself will not be needed in what follows. It appears to be a feature of second descents that the explicit construction of models requires far less work.

THEOREM 6.4. *Given a ‘Black Box’ as above, a Weierstrass equation for E and some element in R^\times representing an n -covering (C, ρ) of E with trivial obstruction, we can explicitly compute a set of defining equations for the image of the Brauer-Severi diagram $C \rightarrow \mathbb{P}^{n-1}$ corresponding to (C, ρ) . This produces a genus one normal curve of degree n defined by:*

- an equation $y^2 = f(x)$ with $f \in K[x]$ separable of degree 4, when $n = 2$;
- a ternary cubic form with coefficients in K , when $n = 3$;
- a set of $n(n - 3)/2$ linearly independent quadrics over K when $n \geq 4$.

CHAPTER II

The Descent Map

Let p be a prime, K a perfect field of characteristic not equal to p , E an elliptic curve over K and $\rho : C \rightarrow E$ a p -covering of E . In this chapter we study the set $\text{Cov}_0^{(p)}(C/K)$ of isomorphism classes of p -coverings of C with trivial obstruction. The main tool here is the *descent map* (defined in section 3). Much like the situation for p -descents on elliptic curves, this allows us to embed $\text{Cov}_0^{(p)}(C/K)$ into a quotient of some étale algebra. When K is a number field, this set contains the p -Selmer set so this becomes the theoretical foundation for performing second p -descents. For an outline of the contents of this (and the next) chapter we refer the reader back to the introduction.

Throughout this chapter we make the following assumptions on C .

- $\text{Pic}(C) = \text{Pic}_K(C)$, i.e. every K -rational divisor class can be represented by some K -rational divisor.
- $\text{Cov}^{(p)}(C/K) \neq \emptyset$, i.e. there exists a p -covering of C defined over K .

REMARK: These assumptions are satisfied when K is a number field and C is everywhere locally solvable. The first is a result of Cassels [Ca2, Theorem 1.2]; it is a consequence of the local-global principle for the Brauer group of K . The second is a result of Tate (appearing in the same article of Cassels, lemma 6.1). It is ultimately a consequence of the local-global principle for $H^1(K, E[p])$. These results were also mentioned in the previous chapter.

From the second assumption above it follows that (C, ρ) has trivial obstruction. The Brauer-Severi diagram corresponding to (C, ρ) gives a model for C as a genus one normal curve of degree p in \mathbb{P}^{p-1} . We fix defining equations of the following form. For $p = 2$, C is a double cover of the projective line ramified in four points. We have a model in the $(1, 1, 2)$ -weighted projective plane given by an equation $u_3^2 = c \cdot f(u_1, u_2)$, where $f(u_1, u_2) \in K[u_1, u_2]$ is a binary quartic monic in u_1 and $c \in K^\times$. The typical affine model is given by setting $u_2 = 1$. For $p = 3$, $C \subset \mathbb{P}^2$ is defined by the vanishing of some ternary cubic form $U(u_1, u_2, u_3) \in K[u_1, u_2, u_3]$. For larger p , the model is as a (noncomplete) intersection of $p(p-3)/2$ quadrics $Q_i(u_1, \dots, u_p) \in K[u_1, \dots, u_p]$.

Note that the model for C determines the class of (C, ρ) in $H^1(K, E[p])$ up to an automorphism of E as an elliptic curve. The model together with the structure of C as a torsor under E determines a unique class in $H^1(K, E[p])$. The torsor structure is given by fixing an isomorphism $\psi : C \rightarrow E$, defined over \bar{K} , such that $\rho = p \circ \psi$. The set $\text{Cov}_0^{(p)}(C/K)$ studied in this chapter does not, however, depend on this extra data. In particular, the results of this chapter apply to any genus one normal curve of degree p satisfying the two assumptions above.

1. The fake Selmer set

To motivate the coming material, we give here a summary of a naive attempt at a second p -descent. Those who are familiar with descent on elliptic curves and other objects would probably consider this to be the obvious thing to try. We do not give complete proofs here since the purpose is primarily motivational.

DEFINITION 1.1. *When $p \geq 3$, we say that a point $x \in C$ is a flex point if there is a hyperplane in \mathbb{P}^{p-1} meeting C in x with multiplicity p . For $p = 2$, we define the flex points to be the ramification points of the double cover of \mathbb{P}^1 . In both cases we denote the set of flex points of C by X .*

REMARK: The terminology generalizes the classical notion of flex points of plane cubics.

Since the flex points are defined by a geometric property, they are stable under the action of G_K . In other words, X is a G_K -set. The action of E on C restricts to an action of $E[p]$ on X . One can see this by noting that the flex points are precisely the points of C lying above 0_E under the covering map. Thus X is an $E[p]$ -torsor; its class in $H^1(K, E[p])$ is the same as that of the p -covering (C, ρ) (see [**CFOSS-I**, Section 1]). It follows also that $\#X = p^2$. We use F (for ‘flex algebra’) to denote the corresponding étale K -algebra,

$$F := \text{Map}_K(X, \bar{K}),$$

and use $[\mathbf{x}]$ to denote the map $X \rightarrow \text{Div}(\bar{C})$ whose value at $x \in X$ is the divisor $[x]$.

To perform a p -descent on an elliptic curve, one uses a Galois equivariant family of functions with zeros of order p at the nontrivial p -torsion points. For performing descent on C , the analog is a family of functions with zeros of order p at the flex points of C . To obtain such a function we choose a linear form $\tilde{t} \in F[u_1, \dots, u_p]$ whose divisor $\text{div}(\tilde{t}) \in \text{Div}(C \otimes_K F) = \text{Map}_K(X, \text{Div}(\bar{C}))$ is equal to $p[\mathbf{x}]$. Existence follows from the definition of a flex point. For $p = 2$ we take the linear form $u_1 - \theta u_2$, where θ denotes the map sending a flex point to its u_1 -coordinate in the affine model given by setting $u_2 = 1$ (cf. the examples following Proposition I.3.1). We then choose some linear form $u \in K[u_1, \dots, u_p]$ which cuts out a divisor on C that is disjoint from X . Their ratio gives a rational function $t := \tilde{t}/u \in \kappa(C \otimes_K F)^\times$. For $p = 2$, choosing $u = u_2$ recovers the function denoted ‘ $u_1 - \theta$ ’ in the examples of I.3.

The divisor of t is $\text{div}(t) = p[\mathbf{x}] - \text{div}(u)$. We find ourselves in the situation of Proposition I.3.1, which yields a homomorphism

$$\Phi_{fake} : \text{Pic}_K(C) \longrightarrow \frac{F^\times}{K^\times F^{\times p}}.$$

Recall that for a divisor class $\Xi \in \text{Pic}_K(C)$, represented by a K -rational divisor $d = \sum_P n_P [P]$ with support disjoint from X and the zeros of u , $\Phi_{fake}(\Xi)$ is equal to the class of $\prod t(P)^{n_P}$ modulo $K^\times F^{\times p}$. In particular, if $P \in C(K) = \text{Pic}_K^1(C)$ is any point which is neither a flex nor a pole of t , then its image under this map is given by evaluating \tilde{t} at some set of coordinates for P in K .

Suppose (D, π) is a p -covering of C with trivial obstruction. Then by lemma I.6.3 there is a model for D in $\mathbb{P}^{p^2-1}(z_1 : \dots : z_{p^2})$ with the property that the pull-back of any flex is a hyperplane section. This means we can choose a Galois equivariant family of linear forms $h_x \in \bar{K}[z_1, \dots, z_{p^2}]$, indexed by $x \in X$, such that the divisor on D cut

out by h_x is equal to $\pi^*[x]$. The Galois equivariance means $h_x^\sigma = h_{x^\sigma}$ for all $\sigma \in G_K$. So $h = (h_x)$ can be considered as a linear form with coefficients in F cutting out the divisor $\pi^*[\mathbf{x}]$ on $D \otimes_K F$.

On the other hand, the divisor of \tilde{t} is $p[\mathbf{x}]$. So there exists some $\delta \in F^\times$ such that the relation $\tilde{t} \circ \pi = \delta h^p$ holds in the coordinate ring of $D \otimes_K F$. We use this to define a map

$$\tilde{\Phi}_{fake} : \text{Cov}_0^{(p)}(C/K) \rightarrow F^\times / K^\times F^{\times p},$$

sending (D, π) to the class of δ . Of course one must check that this is well-defined; for the sake of this discussion we will just assume it (cf. theorem 3.2). We will refer to this as the *fake descent map*. This is simply to distinguish it from the map we define with theorem 3.2 below.

The definition is functorial in the base field and this fake descent map has the property that if $K \subset L$ is any extension and $Q \in D(L)$ is an L -rational point, then

$$\tilde{\Phi}_{fake}((D, \pi)) \equiv \Phi_{fake}(\pi(Q)) \pmod{L^\times F_L^{\times p}}.$$

Recall that the p -coverings of C partition its K -rational points. This property says that the map $\Phi_{fake} : C(K) = \text{Pic}_K^1(C) \rightarrow F^\times / K^\times F^{\times p}$ factors through the set of equivalence classes determined by this partition. Thus it can be used to obtain information on the p -coverings of C .

Now specialize to the case that $K = k$ is a number field and use F_v to denote $F \otimes k_v$ for a completion k_v of k . We have a commutative diagram:

$$\begin{array}{ccc} C(k) & \xrightarrow{\Phi_{fake}} & \frac{F^\times}{k^\times F^{\times p}} \\ \downarrow & & \downarrow \prod \text{res}_v \\ \prod_v C(k_v) & \xrightarrow{\prod \Phi_{fake,v}} & \prod_v \frac{F_v^\times}{k_v^\times F_v^{\times p}} \end{array}$$

We make the following definition.

DEFINITION 1.2. *The fake p -Selmer set of C over k is the set*

$$\text{Sel}_{fake}^{(p)}(C/k) = \{ \delta \in F^\times / k^\times F^{\times p} : \text{res}_v(\delta) \in \Phi_{fake,v}(C(k_v)) \text{ for all } v \}.$$

REMARK: For $p = 2$, this coincides with the definitions in [BS, Sta]. Occasionally one also sees the additional condition that $c \cdot N_{F/k}(\delta) \in k^{\times 2}$, where c is the leading coefficient of the binary quartic defining C . While perhaps useful in practice, this is not needed in the definition since any global element that is everywhere locally a square is a square and this condition is satisfied for the local images at all primes.

Suppose $(D, \pi) \in \text{Sel}^{(p)}(C/k)$, and $\tilde{\Phi}_{fake}((D, \pi)) = \delta$. Then for every v , $\text{res}_v(\delta)$ is the image under $\Phi_{fake,v}$ of some k_v -rational point on C . This shows that $\tilde{\Phi}_{fake}$ maps the p -Selmer set to the fake p -Selmer set. In particular, if the fake Selmer set is empty, then so are both $\text{Sel}^{(p)}(C/k)$ and $C(k)$.

When $p = 2$ it is known (we recover a proof in section 6) that the fake descent map is two¹ to one onto its image. Each fiber consists of a pair of coverings which differ

¹If C has a k -rational flex point, then actually the map is injective. This case is however uninteresting; in addition to having an obvious k -rational point, C is trivial as a 2-covering of its Jacobian.

only by the hyperelliptic involution on C . In particular, the two underlying curves are k -isomorphic.

The condition that some δ in the image of $\tilde{\Phi}_{fake}$ lie in $\text{Sel}_{fake}^{(2)}(C/k)$ is that, for each prime v , at least one of the two 2-coverings in the fiber above δ has a k_v -point. However, since the two 2-coverings are k -isomorphic, either both have a k_v -point or neither does. This allows one to conclude that $\tilde{\Phi}_{fake} : \text{Sel}^{(2)}(C/k) \rightarrow \text{Sel}_{fake}^{(2)}(C/k)$ is surjective. Thus, computing the fake 2-Selmer set gives an effective means of performing a second 2-descent.

Those familiar with second 2-descents (or 2-descents on hyperelliptic curves) will recall that the ambiguity ultimately comes down to a choice of square root of $c \cdot N_{F/k}(\delta) \in k^{\times 2}$. For larger p , the fibers can also be seen as parameterizing choices of p -th roots, but now in a certain k -algebra (subject to various conditions and up to a certain equivalence - cf. corollary 5.5). Consequently, the fibers can be larger and less well-behaved. Moreover, the argument above no longer works since the coverings in a given fiber need not be k -isomorphic as curves.

To deal with these issues, we need to somehow tease out the information ignored by this fake descent map. Ultimately, this will require us to use a descent map induced by Galois equivariant families of functions on C whose zero divisors may be supported on multiple points of X , perhaps with higher multiplicities. For example, in the case $p = 2$, the ambiguity can be eliminated by using the additional function $u_3/u_2^2 \in \kappa(\bar{C})^\times$, whose zero divisor is the sum of the four flex points. In practice this is hardly necessary, but it is indicative of the situation for odd p .

2. The linear part of the descent map

Throughout this section we work over K , an arbitrary field of characteristic not equal to p , keeping the notation and assumptions laid out above. In this section we suppress covering maps from the notation: when we write $D \in \text{Cov}^{(p)}(C/K)$ it is implicit that this means the class of some covering $\pi : D \rightarrow C$.

Since we have assumed that there exists a p -covering of C , the set $\text{Cov}^{(p)}(C/K)$ is a principal homogeneous space for $H^1(k, E[p])$ (cf. Definition I.5.1). The action of a class represented by a cocycle ξ on a covering D is given by twisting. We use D_ξ to denote the twist of D by ξ . Both D and ξ have canonical images in $H^1(k, E[p^2])$ and the action of twisting coincides with the group law there. Namely, the image of D_ξ is the sum of the images of D and ξ . The obstruction map Ob_n was defined in section I.6. We now identify how this changes under the action of twisting.

LEMMA 2.1. *For $D \in \text{Cov}^{(p)}(C/K)$ and $\xi \in H^1(K, E[p])$ we have*

$$\text{Ob}_{p^2}(D_\xi) = C \cup_p \xi + \text{Ob}_{p^2}(D),$$

where \cup_n denotes the cup product associated to the Weil pairing of level n .

REMARK: Rewriting this as $C \cup_p \xi = \text{Ob}_{p^2}(D + \xi) - \text{Ob}_{p^2}(D)$, we can interpret the cup product $C \cup_p \xi$ as the directional derivative of Ob_{p^2} along ξ .

PROOF: For the proof, we identify D , D_ξ and ξ with their images in $H^1(K, E[p^2])$. We know that Ob_n is quadratic, and that the associated bilinear form is given by \cup_n (see

I.6). This means that

$$D \cup_{p^2} \xi = \text{Ob}_{p^2}(D_\xi) - \text{Ob}_{p^2}(D) - \text{Ob}_{p^2}(\xi).$$

The compatibility of the obstruction maps of different levels (lemma I.6.2) shows that $\text{Ob}_{p^2}(\xi) = 0$. On the other hand, the Weil pairings of levels mn and n satisfy the compatibility condition (see [Sil, III.8]):

$$\text{for all } S \in E[mn] \text{ and } T \in E[m], e_{mn}(S, T) = e_m(nS, T).$$

For the cup product on the left-hand side above this means

$$D \cup_{p^2} \xi = (p_*D) \cup_p \xi = C \cup_p \xi,$$

which completes the proof. \square

We use C^\perp to denote the annihilator of C with respect to \cup_p , i.e.

$$C^\perp = \{\xi \in H^1(K, E[p]) : C \cup_p \xi = 0\}.$$

COROLLARY 2.2. *The set $\text{Cov}_0^{(p)}(C/K)$ is either empty or is a principal homogeneous space for $C^\perp \subset H^1(K, E[p])$.*

PROOF: This is clear from the lemma and the fact that the action of $H^1(K, E[p])$ on $\text{Cov}^{(p)}(C/K)$ is compatible with the group law in $H^1(K, E[p^2])$ \square

Affine Maps. As noted above, the set X of flex points has the structure of a K -torsor under $E[p]$. As such, X may be identified with the affine space underlying the 2-dimensional \mathbb{F}_p -vector space $E[p]$. The action of G_K on X factors through the affine general linear group, which is an extension of the general linear group by the group of translations:

$$1 \rightarrow E[p] \rightarrow \text{AGL}(X) \rightarrow \text{GL}(E[p]) \rightarrow 1.$$

Here $E[p]$ acts on X by translations and $\text{GL}(E[p])$ acts on $E[p]$ in the obvious way.

In general, if V, W are vector spaces and \mathbb{A} denotes the affine space underlying V , then a map $\phi : \mathbb{A} \rightarrow W$ is said to be affine if, for all $x \in \mathbb{A}$ and $P, Q \in V$, one has

$$\phi(x + P + Q) + \phi(x) = \phi(x + P) + \phi(x + Q).$$

Geometrically, this says that the sums of the values of ϕ on the two pairs of opposite vertices of any parallelogram in \mathbb{A} are equal. We define $\text{Aff}(\mathbb{A}, W)$ to be the vector space of affine maps from \mathbb{A} to W .

Given an affine map $\phi : \mathbb{A} \rightarrow W$ and $x \in \mathbb{A}$, we can obtain a linear map $\Lambda_{\phi, x} : V \rightarrow W$ by ‘projecting onto the linear part’. This is defined by $\Lambda_{\phi, x}(P) = \phi(x + P) - \phi(x)$.

LEMMA 2.3. *$\Lambda_{\phi, x}$ is linear and does not depend on the choice for x .*

PROOF: First we show that $\Lambda_{\phi, x}$ is independent of x . For this let $x' \in \mathbb{A}$ be any other point. There is a uniquely determined $Q \in V$ such that $x' = x + Q$. Then for any $P \in V$,

$$\Lambda_{\phi, x}(P) - \Lambda_{\phi, x'}(P) = \phi(x + P) - \phi(x) - \phi(x + P + Q) + \phi(x + Q),$$

which is equal to 0, since ϕ is affine. To show $\Lambda_\phi = \Lambda_{\phi,x}$ is linear, let $P' \in V$. Then using that ϕ is affine

$$\begin{aligned}\Lambda_\phi(P + P') &= \phi(x + P + P') - \phi(x) \\ &= \phi(x + P) + \phi(x + P') - 2\phi(x) \\ &= \Lambda_\phi(P) + \Lambda_\phi(P').\end{aligned}$$

□

This projection gives rise to a surjective linear map $\text{Aff}(\mathbb{A}, W) \rightarrow \text{Hom}(V, W)$. One can easily verify that the kernel is the space of constant maps. Thus we have an exact sequence

$$0 \rightarrow W \rightarrow \text{Aff}(\mathbb{A}, W) \rightarrow \text{Hom}(V, W) \rightarrow 0.$$

Now return to the case $V = E[p]$ and $\mathbb{A} = X$. We consider μ_p as an \mathbb{F}_p -vector space written multiplicatively. It naturally embeds in $\text{Aff}(X, \mu_p)$ as the subspace of constant maps. The group $\text{Aff}(X, \mu_p)$ itself may be identified with a subgroup of $\text{Map}(X, \bar{K})$. As such it inherits a natural action of G_K . We have a short exact sequence of G_K -modules

$$(2.1) \quad 1 \rightarrow \mu_p \rightarrow \text{Aff}(X, \mu_p) \rightarrow \text{Hom}(E[p], \mu_p) \rightarrow 0.$$

The G_K -module $E[p]$ is self-dual via the Weil pairing. Namely, we can identify $E[p]$ with $\text{Hom}(E[p], \mu_p)$ via

$$E[p] \ni P \mapsto e_p(P, -) \in \text{Hom}(E[p], \mu_p).$$

REMARK: Alternatively, one can make this identification using $P \leftrightarrow e_p(-, P)$. Since the Weil pairing is alternating, the two differ by a sign. This controls the factor of -1 in the next lemma. We have made our choice in deference to the formulation of proposition 4.1 below.

Making this identification in the exact sequence (2.1) above and taking Galois cohomology we obtain an exact sequence

$$(2.2) \quad H^1(K, \mu_p) \rightarrow H^1(K, \text{Aff}(X, \mu_p)) \rightarrow H^1(K, E[p]) \xrightarrow{\Upsilon} \text{Br}(K)[p].$$

Here we have also identified $H^2(K, \mu_p)$ with the p -torsion in the Brauer group of K . The next lemma identifies C^\perp with the kernel of Υ .

LEMMA 2.4. $\Upsilon(\xi) = -C \cup_p \xi$.

REMARK: We may consider $\text{Cov}^{(p)}(C/K)$ as the affine space underlying the \mathbb{F}_p -vector space $H^1(K, E[p])$. With this interpretation the obstruction map $\text{Ob}_{p^2} : \text{Cov}^{(p)}(C/K) \rightarrow \text{Br}(K)[p]$ is affine, as one can see from lemma 2.1. Lemma 2.4 identifies Υ (up to sign) as the corresponding linear map obtained by projecting.

PROOF: Recall that $\rho : C \rightarrow E$ denotes the covering map and that $\psi : C \rightarrow E$ is an isomorphism (defined over some extension of K) such that $p \circ \psi = \rho$. For any $\sigma \in G_K$, the map $\psi^\sigma - \psi$ corresponds to translation by an element of $E[p]$. This defines a cocycle representing the class of C in $H^1(K, E[p])$. The cup product $-C \cup_p \xi$ is the class of the 2-cocycle

$$G_K \times G_K \ni (\sigma, \tau) \mapsto e_p(\psi - \psi^\tau, \xi_\sigma^\tau) \in \mu_p.$$

Now let $\xi \in H^1(K, E[p])$. Υ is a connecting homomorphism, so to compute $\Upsilon(\xi)$ we first choose a lift of ξ to a cochain with values in $\text{Aff}(X, \mu_p)$. For any $P \in E[p]$, we claim that the map $\phi_P : X \ni x \mapsto e_p(P, \psi(x)) \in \mu_p$ is affine and that its image under $\text{Aff}(X, \mu_p) \rightarrow E[p]$ is P . To see that it is affine, let $x \in X$ and $Q, R \in E[p]$. Using bilinearity of the Weil pairing we have

$$\begin{aligned} \phi_P(x + Q + R) \cdot \phi_P(x) &= e_p(P, \psi(x + Q + R)) \cdot e_p(P, \psi(x)) \\ &= e_p(P, \psi(x) + Q + R) \cdot e_p(P, \psi(x)) \\ &= e_p(P, \psi(x) + Q) \cdot e_p(P, \psi(x) + R) \\ &= \phi_P(x + Q) \cdot \phi_P(x + R). \end{aligned}$$

The image of ϕ_P in $\text{Hom}(E[p], \mu_p)$ is given by projecting onto the linear part. This is the map

$$R \mapsto \phi_P(x + R)/\phi_P(x) = e_p(P, \psi(x) + R)/e_p(P, \psi(x)) = e_p(P, R).$$

The identification of $E[p]$ with $\text{Hom}(E[p], \mu_p)$ is given by

$$E[p] \ni P \leftrightarrow e_p(P, -) \in \text{Hom}(E[p], \mu_p),$$

so the image of ϕ_P in $E[p]$ is P .

Thus $\Upsilon(\xi)$ is given by the coboundary of the cochain $\sigma \mapsto e_p(\xi_\sigma, \psi) = e_p(-\psi, \xi_\sigma) \in \text{Aff}(X, \mu_p)$. Here $e_p(-\psi, \xi_\sigma)$ is the map $x \mapsto e_p(-\psi(x), \xi_\sigma)$. The value of the coboundary on a pair $(\sigma, \tau) \in G_K \times G_K$ is given by

$$\frac{e_p(-\psi, \xi_\sigma)^\tau \cdot e_p(-\psi, \xi_\tau)}{e_p(-\psi, \xi_{\sigma\tau})} = \frac{e_p(-\psi^\tau, \xi_\sigma^\tau)}{e_p(-\psi, \xi_\sigma^\tau)} = e_p(\psi - \psi^\tau, \xi_\sigma^\tau).$$

This is the same as the cup product computed above, so the lemma is proven. \square

When is $\text{Cov}_0^{(p)}(C/k)$ nonempty? The material of this subsection will not be needed in what follows. Suppose C is defined over a number field k and is everywhere locally solvable. It is natural to ask if $\text{Cov}_0^{(p)}(C/k)$ is always nonempty. If so, then every element of $\text{III}(E/k)[p]$ lifts (under multiplication by p) to an element in $H^1(k, E)[p^2]$ of index dividing p^2 (cf. the remark following lemma I.6.2). We offer the following answer in the case $p = 2$.

THEOREM 2.5. *If $C \in \text{Sel}^{(2)}(E/k)$, then $\text{Cov}_0^{(2)}(C/k) \neq \emptyset$.*

PROOF: To begin with let p be an arbitrary prime. Since we have assumed C to have points everywhere locally we have $\text{Ob}_p(C) = 0$. From this and the compatibility of Ob_p and Ob_{p^2} (see I.6.2) it follows that the obstruction algebra associated to any $D \in \text{Cov}^{(p)}(C/k)$ is actually p -torsion. So from lemma 2.4 it follows that the image of $\text{Ob}_{p^2} : \text{Cov}^{(p)}(C/k) \rightarrow \text{Br}(k)[p]$ is a coset of the image of Υ . Evidently, $\text{Cov}_0^{(p)}(C/k) \neq \emptyset$ if and only if this coset is equal to the image of Υ .

Extending the exact sequence (2.2) defining Υ we have

$$H^1(k, \text{Aff}(X, \mu_p)) \rightarrow H^1(k, E[p]) \xrightarrow{\Upsilon} \text{Br}(k)[p] \xrightarrow{\alpha} H^2(k, \text{Aff}(X, \mu_p)).$$

By exactness, the image of Υ is the kernel of α . The discussion above shows that the composition $\alpha \circ \text{Ob}_{p^2} : \text{Cov}^{(p)}(C/k) \rightarrow H^2(k, \text{Aff}(X, \mu_p))$ is constant, equal to say $a \in H^2(k, \text{Aff}(X, \mu_p))$. Moreover, $\text{Cov}_0^{(p)}(C/k) \neq \emptyset$ if and only if $a = 0$. For any prime v , we also see that $\text{Cov}_0^{(p)}(C/k_v) \neq \emptyset$ if and only if $\text{res}_v(a) = 0$.

On the other hand, for each prime v , there exists some $D_v \in \text{Cov}^{(p)}(C/k_v)$ with $D_v(k_v) \neq \emptyset$. Such a covering must have trivial obstruction (over k_v), so everywhere locally $\text{Cov}_0^{(p)}(C/k_v) \neq \emptyset$. Consequently $a \in \text{H}^2(k, \text{Aff}(X, \mu_p))$ is everywhere locally trivial. To prove the proposition it is enough to show that $\text{H}^2(k, \text{Aff}(X, \mu_2))$ satisfies the Hasse principle, i.e. any class in $\text{H}^2(k, \text{Aff}(X, \mu_2))$ that is everywhere locally trivial is trivial.

For the moment we will continue to work with an arbitrary prime p . To ease notation let $M = \text{Aff}(X, \mu_p)$. By Poitou-Tate duality [CoN, 8.6.7], the Hasse principle holds or fails simultaneously for $\text{H}^2(k, M)$ and $\text{H}^1(k, M^\vee)$, where $M^\vee = \text{Hom}(M, \mu_p)$. Our strategy is to write down a map from $\text{H}^1(k, M^\vee)$ to a group known to satisfy the Hasse principle. The kernel of this map is finite and depends only on the action of G_k on the flex points of C . For fixed p there are only finitely many possibilities for the action, and for each, the kernel can be computed. For $p = 2$ it turns out that the map is always injective.

Let $R = \text{Map}_k(M, \bar{k})$ be the algebra of all G_k -equivariant maps from M to \bar{k} and let Q denote the quotient of $\mu_p(\bar{R}) = \text{Map}(M, \mu_p)$ by the subspace consisting of maps that are homomorphisms. Then we have an exact sequence

$$0 \rightarrow M^\vee \rightarrow \mu_p(\bar{R}) \xrightarrow{q} Q \rightarrow 0.$$

Taking the Galois cohomology of this sequence over k and its completions we obtain a diagram with exact rows

$$\begin{array}{ccccccc} \mu_p(R) & \xrightarrow{q} & \text{H}^0(k, Q) & \longrightarrow & \text{H}^1(k, M^\vee) & \longrightarrow & \text{H}^1(k, \mu_p(\bar{R})) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \prod_v \mu_p(R_v) & \xrightarrow{\prod q_v} & \prod_v \text{H}^0(k_v, Q) & \longrightarrow & \prod_v \text{H}^1(k_v, M^\vee) & \longrightarrow & \prod_v \text{H}^1(k_v, \mu_p(\bar{R})) \end{array}$$

The Grunwald-Wang theorem [CoN, 9.1.11], implies that the rightmost vertical map is injective. So if $\text{H}^1(k, M^\vee) \rightarrow \text{H}^1(k, \mu_p(\bar{R}))$ is injective, then the Hasse principle holds for $\text{H}^1(k, M^\vee)$.

The action of G_k on X factors through the affine general linear group $\text{AGL}_2(\mathbb{F}_p)$ and determines the action on μ_p . The actions of G_k on X , μ_p and any modules derived from the two (e.g. M^\vee and $\mu_p(\bar{R})$) depend only on its image in $\text{AGL}_2(\mathbb{F}_p)$. If this is denoted by \mathcal{G} , then we have a commutative diagram with exact rows:

$$\begin{array}{ccccccc} \text{H}^0(k, \mu_p(\bar{R})) & \xrightarrow{q} & \text{H}^0(k, Q) & \longrightarrow & \text{H}^1(k, M^\vee) & \longrightarrow & \text{H}^1(k, \mu_p(\bar{R})) \\ \parallel & & \parallel & & & & \\ \text{H}^0(\mathcal{G}, \mu_p(\bar{R})) & \xrightarrow{q} & \text{H}^0(\mathcal{G}, Q) & \longrightarrow & \text{H}^1(\mathcal{G}, M^\vee) & \longrightarrow & \text{H}^1(\mathcal{G}, \mu_p(\bar{R})) \end{array}$$

The kernel of the map $\text{H}^1(k, M^\vee) \rightarrow \text{H}^1(k, \mu_p(\bar{R}))$ is thus isomorphic to the finite group $\text{H}^0(\mathcal{G}, Q)/q(\text{H}^0(\mathcal{G}, \mu_p(\bar{R})))$. Note also that this only depends on the conjugacy class of \mathcal{G} in $\text{AGL}_2(\mathbb{F}_p)$. For a subgroup $H \subset \text{AGL}_2(\mathbb{F}_p)$, let us use $\mathcal{R}(H)$ to denote $\text{H}^0(H, Q)/q(\text{H}^0(H, \mu_p(\bar{R})))$.

The proof of the proposition is thus reduced to verification of the following statement:

$$\text{For any subgroup } H \subset \text{AGL}_2(\mathbb{F}_2) \simeq S_4 \text{ we have } \mathcal{R}(H) = 0.$$

Up to conjugacy there are 11 subgroups of S_4 . For each H , $\mathcal{R}(H)$ is a combinatorial object that can be computed using linear algebra over \mathbb{F}_2 . We have performed this computation in MAGMA and verified, for each H , that $\mathcal{R}(H) = 0$. \square

For odd p , this Hasse principle can fail - we have the following example. The cubic curve

$$C : x^3 + x^2z + 5xy^2 - 4xyz - 9xz^2 + 2y^3 - 2y^2z + 9yz^2 - 6z^3 = 0$$

defined over \mathbb{Q} is everywhere locally solvable. If F' denotes the extension of \mathbb{Q} obtained by adjoining the coordinates of all flex points of C , then its Galois group $G = \text{Gal}(F'|\mathbb{Q})$ is an extension of $\mathbb{Z}/2\mathbb{Z}$ by $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$. The subgroup $G' = \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ acts transitively on X and the quotient corresponds to adjoining the cube roots of unity to \mathbb{Q} . A computation as in the lemma shows that $\mathcal{R}(G) \simeq \mathcal{R}(G') \neq 0$ and that up to conjugacy these are the only subgroups of $\text{AGL}_2(\mathbb{F}_3)$ with this property. Above any prime v of \mathbb{Q} there are at least 3 primes of F' (only the ramified primes need to be checked), so all decomposition groups are of index ≥ 3 in G . In particular, none of them are isomorphic to G or G' . So, for each v , $\mathcal{R}(G_v) = 0$ and the Hasse principle must fail. Indeed, the kernels of $\mathcal{R}(G) \rightarrow \prod_v \mathcal{R}(G_v)$ and $H^1(\mathbb{Q}, \text{Aff}(X, \mu_3)^\vee) \rightarrow \prod H^1(\mathbb{Q}_v, \text{Aff}(X, \mu_3)^\vee)$ are isomorphic.

Of course this does not tell us that there are no 3-coverings of C with trivial obstruction and, for the curve in question, this is not the case. With corollary 5.4 below we give a concrete algebraic realization of $\text{Cov}_0^{(3)}(C/k)$ as a subquotient of some étale k -algebra defined by certain norm conditions. For the curve above one can check that these norm conditions are satisfiable and so $\text{Cov}_0^{(3)}(C/k) \neq \emptyset$. Our gut feeling is that Theorem 2.5 should fail for odd p , but it seems very difficult to produce counter-examples.

Making cohomology groups explicit. We have identified C^\perp with the kernel of Υ . By exactness of the sequence (2.2) defining Υ , this is the same as the image of $H^1(K, \text{Aff}(X, \mu_p))$ in $H^1(K, E[p])$. This suggests that we should look for a practical description of $H^1(K, \text{Aff}(X, \mu_p))$.

Recall that F denotes the flex algebra $\text{Map}_K(X, \bar{K})$ and that $\mu_p(\bar{F}) = \text{Map}(X, \mu_p)$. We have a canonical monomorphism $\text{Aff}(X, \mu_p) \hookrightarrow \mu_p(\bar{F})$; this is simply the observation that an affine map is a map. To obtain a description of $H^1(K, \text{Aff}(X, \mu_p))$ we want to extend this to a short exact sequence and take its Galois cohomology. To make this useful, we need a better description of the quotient.

The case $p = 2$. A map is affine if the products of its values on either pair of opposite vertices of a parallelogram in X are the same. For $p = 2$, there is only one (nondegenerate) parallelogram in X . The norm map $N_{F/K} : F \rightarrow K$ takes $\phi \in \bar{F} = \text{Map}(X, \bar{K})$ to the product of its values on all vertices of this parallelogram. It is then easy to see that a map $\phi \in \bar{F}^\times = \text{Map}(X, \bar{K}^\times)$ is in $\text{Aff}(X, \mu_2)$ if and only if $\phi^2 = N_{F/K}(\phi) = 1$. To encode this, we define a map

$$\partial : F \ni \phi \mapsto (\phi^2, N_{F/K}(\phi)) \in F \times K.$$

LEMMA 2.6. *We have $(\partial\bar{F}^\times)^{G_K} = \{(\delta, \varepsilon) \in F^\times \times K^\times \mid N_{F/K}(\delta) = \varepsilon^2\}$ and*

$$H^1(K, \text{Aff}(X, \mu_2)) \simeq \frac{(\partial\bar{F}^\times)^{G_K}}{\partial F^\times}.$$

PROOF: It was noted above that the kernel of ∂ on \bar{F}^\times is the group of affine maps to μ_2 . This gives an exact sequence

$$1 \rightarrow \text{Aff}(X, \mu_2) \rightarrow \bar{F}^\times \xrightarrow{\partial} \partial\bar{F}^\times \rightarrow 1.$$

Its long exact sequence, together with Hilbert's Theorem 90, gives

$$F^\times \xrightarrow{\partial} (\partial\bar{F}^\times)^{G_K} \rightarrow H^1(K, \text{Aff}(X, \mu_2)) \rightarrow H^1(K, \bar{F}^\times) = 0,$$

from which the second statement follows. The description of $(\partial\bar{F}^\times)^{G_K}$ is clear. \square

REMARK: One should think of $(\partial\bar{F}^\times)^{G_K}$ as the set of elements in F^\times which have square norm, together with a choice of square root for this norm.

REMARK: Let $Y \subset \text{Div}(\bar{C})$ be the set containing the divisor $\sum_{x \in X} [x]$ and the divisors, $2[x]$ for $x \in X$. Then Y is a G_K -set derived from X . Its corresponding étale K -algebra is isomorphic to $F \times K$ and $\partial : F \rightarrow F \times K$ is simply the 'induced norm map' defined in section I.3.

The case $p \geq 3$. Now assume p is odd. We want to characterize the affine maps as above. One way would be to write down a map which encodes the property defining affine maps. This would lead to consideration of a K -algebra of degree $O(p^6)$. We would like to get away with less. The following easy lemma does this in the case $p = 3$.

LEMMA 2.7. *A map $\phi : \mathbb{A}_{\mathbb{F}_3}^2 \rightarrow \mathbb{F}_3$ is affine if and only if $\sum_{x \in \ell} \phi(x) = 0$ for every affine line $\ell \subset \mathbb{A}_{\mathbb{F}_3}^2$.*

PROOF: Easy. This also follows from lemma 2.8 below. \square

REMARK: When $p > 3$ there are maps with this property that fail to be affine. This is because the map obtained by a projection as in lemma 2.3 may fail to be homogeneous of degree one (cf. lemma I.4.2).

The affine lines in $\mathbb{A}_{\mathbb{F}_3}^2$ correspond to the 12 unordered triples of points of the form

$$\{x, x + P, x - P\},$$

where $x \in \mathbb{A}_{\mathbb{F}_3}^2$, $0 \neq P \in \mathbb{F}_3^2$ and $x \pm P$ denotes the translate of x by $\pm P$. So the property in the lemma can be rewritten as

$$\phi(x + P) + \phi(x - P) = -\phi(x) = 2\phi(x), \text{ for all } x \in \mathbb{A}_{\mathbb{F}_3}^2 \text{ and } P \in \mathbb{F}_3^2 \setminus \{0\}.$$

Recall that a map is affine if the sums of its values on the two pairs of opposite vertices of any parallelogram are equal. The condition above expresses this for the degenerate parallelograms where one pair of opposite vertices coincide.

LEMMA 2.8. *Let p be an odd prime and $\phi : \mathbb{A}_{\mathbb{F}_p}^2 \rightarrow \mathbb{F}_p$. In order that ϕ be affine it is necessary and sufficient that*

$$\phi(x + P) + \phi(x - P) = 2\phi(x),$$

for all $x \in \mathbb{A}_{\mathbb{F}_p}^2$ and $P \in \mathbb{F}_p^2 \setminus \{0\}$.

PROOF: The necessity of this condition is clear from the discussion above. To show sufficiency, assume the condition is satisfied and let $S, T \in \mathbb{F}_p^2$ and $x \in \mathbb{A}_{\mathbb{F}_p}^2$. We must show that $\phi(x + S + T) + \phi(x) = \phi(x + S) + \phi(x + T)$.

Using the condition on the pairs $(x + S, T)$, $(x + T, S)$, $(x, S - T) \in \mathbb{A}_{\mathbb{F}_p}^2 \times \mathbb{F}_p^2$ we get

$$\begin{aligned} \phi(x + S + T) &= 2\phi(x + S) - \phi(x + S - T) \\ \phi(x + S + T) &= 2\phi(x + T) - \phi(x + T - S) \\ 2\phi(x) &= \phi(x + (S - T)) + \phi(x - (S - T)) \end{aligned}$$

Summing these yields,

$$2\phi(x + S + T) + 2\phi(x) = 2\phi(x + S) + 2\phi(x + T),$$

which gives the result since 2 is invertible, as p is assumed to be odd. \square

In order to cut out the affine maps $\text{Aff}(X, \mu_p) \subset \mu_p(\bar{F})$ we would like to write down a ‘norm map’ that encodes this condition. To set up for an eventual application of proposition I.3.1 we want to describe this in terms of divisors on C .

LEMMA 2.9. *The set of divisors of the form $(p - 2)[x] + [x + P] + [x - P] \in \text{Div}(\bar{C})$, with $x \in X$ and $P \in E[p]$, is a G_K -stable set of hyperplane sections of C .*

PROOF: The fact that these divisors form a G_K -set is obvious, since both the flex points of C and the p -torsion points of E are themselves G_K -sets. To see that they are hyperplane sections, we may work geometrically, considering C as an elliptic curve with some flex $x_0 \in X$ as distinguished point. Note that the flex points are then the p -torsion points on the elliptic curve (C, x_0) . Since the model for C is given by the embedding corresponding to the complete linear system $|p[x_0]|$, the hyperplane sections are precisely those divisors linearly equivalent to $p[x_0]$. That the divisors in the lemma are hyperplane sections is then a consequence of the well-known fact that two divisors on an elliptic curve are linearly equivalent if and only if they have the same degree and the same sum. \square

Before proceeding, we fix some notation. We use Y to denote the set of hyperplanes in the lemma. It is a G_K -set and we denote its corresponding étale K -algebra by H (for ‘hyperplane algebra’). Note that Y is derived from X in the sense described in section I.3. As a G_K -set, Y splits as a disjoint union of (at least) two G_K -stable subsets. The first consists of the p^2 hyperplane sections of the form $p[x]$ with $x \in X$. These divisors correspond to pairs (x, P) with $P = 0$. The other consists of those $y \in Y$ associated to some pair (x, P) with $P \neq 0$. These two G_K -subsets will be denoted by Y_1 and Y_2 ; their corresponding étale K -algebras will be denoted by H_1 and H_2 . As G_K -sets, X and Y_1 are isomorphic and so we will identify F with H_1 . Thus H splits as $H \simeq H_1 \times H_2 \simeq F \times H_2$.

For $p = 3$, Y_2 consists of the 12 lines of \mathbb{P}^2 that pass through three distinct flex points of C . For $p \geq 5$, $X \times \frac{E[p] \setminus \{0_E\}}{\{\pm 1\}}$ and Y_2 are isomorphic as G_K -sets, a pair (x, P) corresponding to the hyperplane section $(p - 2)[x] + [x + P] + [x - P]$. From this we see that $\#Y_2 = p^2(p^2 - 1)/2$. There is a canonical projection $Y_2 \ni (x, P) \mapsto x \in X$. Thus, for $p \geq 5$, H_2 may be viewed as an F -algebra of degree $(p^2 - 1)/2$.

Since Y is derived from X as a G_K -set, we have an induced norm map (see section I.3):

$$\partial : F \ni \phi \mapsto \left(y \mapsto \prod_{x \in y} \phi(x) \right) \in H.$$

The product appearing here is to be taken with appropriate multiplicities. So, for example, the value of $\partial(\phi)$ on the divisor $((p-2)[x] + [x+P] + [x-P]) \in Y$ is the product $\phi(x)^{p-2}\phi(x+P)\phi(x-P)$.

Using the characterization of affine maps given by lemma 2.8 we have the following analog of lemma 2.6.

LEMMA 2.10.

$$H^1(K, \text{Aff}(X, \mu_p)) \simeq \frac{(\partial \bar{F}^\times)^{G_K}}{\partial F^\times}.$$

PROOF: By lemma 2.8, a map $\phi \in \bar{F}^\times$ is an affine map to μ_p if and only if the product of its values on any hyperplane section $y \in Y$ is 1 (i.e. if and only if it lies in the kernel of ∂). This gives an exact sequence

$$1 \rightarrow \text{Aff}(X, \mu_p) \rightarrow \bar{F}^\times \xrightarrow{\partial} \partial \bar{F}^\times \rightarrow 1.$$

The result then follows by taking the long exact sequence and using Hilbert's Theorem 90 as in 2.6. \square

REMARK: Under the splitting, $H \simeq F \times H_2$, ∂ splits as $\partial = \partial_1 \times \partial_2$, where the ∂_i are induced by the structure of Y_i as a G_K -set derived from X . Since $\partial_1(\phi) = \phi^p$, we have that

$$(\partial \bar{F}^\times)^{G_K} \subset \{(\delta, \varepsilon) \in F^\times \times H_2^\times : \partial_2(\delta) = \varepsilon^p\}.$$

In the case $p = 2$ this was sufficient to describe $(\partial \bar{F}^\times)^{G_K}$. For odd p equality need not hold in general. One may still, however, think of $(\partial \bar{F}^\times)^{G_K}$ as a subset of elements of F^\times with p -th power norm (in H_2^\times) together with a(n appropriate) choice of p -th root (cf. lemmas 5.3, 5.5).

The notation above can be made compatible with that for $p = 2$ used in the remark following lemma 2.6. Namely, we take Y_1 to be the set of divisors of C that are of the form $2[x]$ for some $x \in X$ and Y_2 to be the set consisting of the divisor $\sum_{x \in X} [x]$. Then $H \simeq H_1 \times H_2 \simeq F \times K$ and the induced norm is the same as the map $\partial : F \rightarrow H$ appearing in lemma 2.6.

Combining the results above, we obtain the following description of C^\perp valid for both even and odd p . Similar statements have appeared in the literature for $p = 2$, most notably [Fi4, Theorem 4.1].

COROLLARY 2.11. *There is an isomorphism $C^\perp \simeq (\partial \bar{F}^\times)^{G_K} / K^\times \partial F^\times$, where for odd p (resp. $p = 2$) we identify K^\times with its image in $H^\times \simeq F^\times \times H_2^\times$ under the diagonal embedding $\iota_p : \alpha \mapsto (\alpha, \alpha)$ (resp. the embedding $\iota_2 : \alpha \mapsto (\alpha, \alpha^2)$).*

PROOF: Lemma 2.4 shows that C^\perp is isomorphic to $H^1(K, \text{Aff}(X, \mu_p))$ modulo the image of $H^1(K, \mu_p)$. Lemma 2.10 (for $p = 2$, use 2.6) and Hilbert's Theorem 90 allow us to identify these groups with $(\partial \bar{F}^\times)^{G_K} / \partial F^\times$ and $K^\times / K^{\times p}$, respectively. We only need to show that the identifications are compatible.

Set $\tilde{p} = \deg(\partial_2)$. Thus $\tilde{p} = p$ for odd p and $\tilde{p} = 2p$ otherwise. Noting that $\bar{K} \subset \bar{F} = \text{Map}(X, \bar{K})$ consists of the constant maps we see that, for $\alpha \in \bar{K}$, $\partial(\alpha) = (\alpha^p, \alpha^{\tilde{p}}) \in \bar{F} \times \bar{H}_2$. For any p , we have the following commutative diagram, where ι_p is the map

in the statement of the lemma.

$$\begin{array}{ccccccccc}
1 & \longrightarrow & \mu_p & \longrightarrow & \bar{K}^\times & \xrightarrow{p} & \bar{K}^\times & \longrightarrow & 1 \\
& & \downarrow & & \downarrow & & \downarrow \iota_p & & \\
1 & \longrightarrow & \text{Aff}(X, \mu_p) & \longrightarrow & \bar{F}^\times & \xrightarrow{\partial} & \partial \bar{F}^\times & \longrightarrow & 1
\end{array}$$

The identifications are made by taking Galois cohomology and using that, by Hilbert's Theorem 90, $H^1(K, -)$ of the middle terms vanish. From this compatibility is clear. \square

From now on, we will always assume K^\times is embedded in H^\times using the map in the lemma.

3. The descent map

Recall (corollary 2.2) that $\text{Cov}_0^{(p)}(C/K)$ is a principal homogeneous space for C^\perp . We have already obtained a more or less concrete description of C^\perp as a subgroup of $H^\times/K^\times\partial F^\times$. The goal now is to give an equally explicit description of $\text{Cov}_0^{(p)}(C/K)$ as a coset of C^\perp inside $H^\times/K^\times\partial F^\times$. This will be achieved by our descent map.

The strategy here will be similar to that used in defining the fake descent map. Recall the Galois equivariant family of functions $t = (t_x) \in \kappa(C \otimes_K F)^\times$ used to define this. The zero divisor of t is the map $p[\mathbf{x}]$, where $[\mathbf{x}] \in \text{Div}(C \otimes_K F) = \text{Map}_K(X, \text{Div}(\bar{C}))$ denotes the map whose value at $x \in X$ is the divisor $[x]$. Similarly we use $[\mathbf{y}]$ to denote the element of $\text{Div}(C \otimes_K H) = \text{Map}_K(Y, \text{Div}(\bar{C}))$ whose value at $y \in Y$ is the divisor $y \in \text{Div}(\bar{C})$. The material of the preceding section indicates that for p -descent on C we want a Galois equivariant family of functions defined over H with zero divisor $[\mathbf{y}]$.

First assume p is odd. By lemma 2.9, the divisors in Y are all hyperplane sections of C . So we can proceed exactly as in section 1. Namely, we choose a linear form $\tilde{\ell} \in H[u_1, \dots, u_p]$ cutting out the divisor $[\mathbf{y}]$. We then choose any linear form $u \in K[u_1, \dots, u_p]$ cutting out a divisor on C that is disjoint from X to obtain a rational function $\ell = \tilde{\ell}/u \in \kappa(C \otimes_K H)^\times$ with

$$\text{div}(\ell) = [\mathbf{y}] - \text{div}(u).$$

For $p = 2$, we use the function denoted $(u_1 - \theta, u_3)$ in the examples of I.3.1. The notation there was in terms of the affine model. In terms of the weighted projective model, we define $\ell \in \kappa(C \otimes_K H)^\times = \text{Map}_K(Y, \kappa(\bar{C})^\times)$ as the map

$$y \mapsto \begin{cases} (u_1 - \theta_x u_2)/u_2 & \text{if } y = 2[x], \text{ and} \\ u_3/u_2^2 & \text{if } y = \sum_{x \in X} [x], \end{cases}$$

where the binary quartic defining C factors as $f(u_1, u_2) = \prod_x (u_1 - \theta_x u_2)$. Note that u_3 has weight 2, so u_3/u_2^2 is indeed a rational function on C . To make the notation compatible with that for odd p , we set $\tilde{\ell} = (u_1 - \theta u_2, u_3)$. One is tempted to refer to $\tilde{\ell}$ as a linear form with coefficients in $H = F \times K$ (even though u_3 is not linear). So as to avoid making constant distinctions between even and odd p we will allow ourselves this abuse of terminology. We also set $u = u_2 \in K[u_1, u_2]$. Recalling that K is embedded in $H = F \times K$ using the map $\alpha \mapsto (\alpha, \alpha^2)$ it makes sense to write $\ell = \tilde{\ell}/u$.

PROPOSITION 3.1. *The function ℓ induces a unique homomorphism*

$$\Phi : \text{Pic}_K(C) \rightarrow H^\times / K^\times \partial F^\times,$$

with the property that the image of any divisor class is given by evaluating ℓ at any K -rational representative with support disjoint from X and $\text{div}(u)$.

PROOF: This follows directly from Proposition I.3.1. \square

REMARK: Under the splitting $Y = Y_1 \amalg Y_2$, we have $H \simeq F \times H_2$ and ℓ corresponds to a pair of rational functions (ℓ_1, ℓ_2) defined over F and H_2 . The divisor of ℓ_1 is $p[\mathbf{x}] - \text{div}(u)$. This means that the fake map Φ_{fake} factors as

$$\Phi_{fake} : \text{Pic}_K(C) \xrightarrow{\Phi} \frac{H^\times}{K^\times \partial F^\times} \xrightarrow{\text{pr}_1} \frac{F^\times}{K^\times F^{\times p}},$$

where pr_1 is the map induced by projection onto the first factor. So any unproven statements appearing in section 1 can be readily deduced from the material presented here by simply ignoring the second factor of H .

Identifying $\text{Pic}_K^1(C)$ with $C(K)$ we can think of Φ as giving a map on the K -points of C . For the points outside X and $\text{div}(u)$, this is simply given by evaluating ℓ . Note also that the homomorphism does not depend on the choice for u . So if we like, we may determine the image of a point by evaluating $\tilde{\ell}$ on some choice of homogeneous coordinates. For this reason we may refer to $\tilde{\ell}$ as the linear form defining the descent map in the following theorem.

THEOREM 3.2. *The choice of linear form $\tilde{\ell}$ determines a unique well-defined map (called the descent map)*

$$\tilde{\Phi} : \text{Cov}_0^{(p)}(C/K) \longrightarrow H^\times / K^\times \partial F^\times$$

with the following property. If $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$ and $K \subset L$ is any extension of fields with $Q \in D(L)$, then

$$\tilde{\Phi}((D, \pi)) \equiv \Phi(\pi(Q)) \pmod{L^\times \partial F_L^\times}.$$

In particular, if $D(K) \neq \emptyset$, then $\tilde{\Phi}((D, \pi))$ is the image of some K -rational point of C under Φ .

REMARK: Recall that $\text{Cov}_0^{(p)}(C/K)$ yields a partition of the K -rational points of C ,

$$C(K) = \coprod_{(D, \pi) \in \text{Cov}_0^{(p)}(C/K)} \pi(D(K)).$$

The defining property says that $\Phi : C(K) \rightarrow H^\times / K^\times \partial F^\times$ is constant on each of the sets appearing in this partition and that the value on each is equal to the image of the corresponding covering under the descent map.

PROOF: Let $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$. By assumption we have a model for (D, π) as a genus one normal curve of degree p^2 in $\mathbb{P}^{p^2-1} = \mathbb{P}^{p^2-1}(z_1 : \dots : z_{p^2})$, where π is defined by homogeneous polynomials, $\pi_i \in K[z_1, \dots, z_{p^2}]$ and the pull-back of any flex point x on C is a hyperplane section \mathfrak{h}_x of D . For any x , \mathfrak{h}_x can be defined by the vanishing of some linear form $h_x \in \bar{K}[z_1, \dots, z_{p^2}]$. Moreover, we can choose these h_x to form a Galois equivariant family. Thus they may be patched together to obtain a linear form $h \in F[z_1, \dots, z_{p^2}]$ cutting out the divisor $\pi^*[\mathbf{x}]$ on $D \otimes_K F$.

Since the zero divisor of ℓ is $[\mathbf{y}] = \partial[\mathbf{x}] \in \text{Div}(C \otimes_K H)$ we see that ∂h and $\tilde{\ell} \circ \pi$ cut out the same divisor on D . It follows that the rational functions $\ell \circ \pi$ and $\partial h/u \circ \pi$ have the same divisor. Hence there exists some $\Delta \in H^\times$ such that

$$(3.1) \quad \ell \circ \pi = \Delta \cdot \left(\frac{\partial h}{u \circ \pi} \right) \text{ in } \kappa(D \otimes_K H)^\times.$$

We define the image of $\tilde{\Phi}((D, \pi))$ to be the class of $\Delta \in H^\times/K^\times \partial F^\times$.

A different choice of forms defining π would change the left-hand side of (3.1) by a factor in K^\times . Similarly, a different choice for the form $h = (h_x)_{x \in X}$ defining $(\mathfrak{h}_x)_{x \in X}$ would change the right-hand side of (3.1) by a factor in ∂F^\times . Thus, having fixed the model for (D, π) in \mathbb{P}^{p^2-1} we get a well-defined element of $H^\times/K^\times \partial F^\times$.

Let us show that this does not depend on the model. Suppose (D', π') is isomorphic to (D, π) , and choose a model for (D', π') in \mathbb{P}^{p^2-1} as genus one normal curve. As above, choose a linear form $h' \in F[z_1, \dots, z_{p^2}]$ cutting out the divisors $\pi'^*[\mathbf{x}]$ on D' . By assumption we have an isomorphism of coverings $\varphi : D' \rightarrow D$ defined over K (i.e. such that $\pi' = \pi \circ \varphi$). Let $\varphi^* : \kappa(D \otimes_K H) \rightarrow \kappa(D' \otimes_K H)$ denote the isomorphism of function fields induced by φ . Applying φ^* to equation (3.1), we obtain a relation in $\kappa(D' \otimes_K H)$,

$$(3.2) \quad \Delta \cdot \left(\frac{\partial(h \circ \varphi)}{u \circ \pi'} \right) = \varphi^* \left(\Delta \cdot \left(\frac{\partial h}{u \circ \pi} \right) \right) = \varphi^*(\ell \circ \pi) = \ell \circ \pi \circ \varphi = \ell \circ \pi'.$$

The divisor on D' cut out by $h \circ \varphi$ is $\pi'^*[\mathbf{x}]$, so the extremal terms in (3.2) define the image of (D', π') under the descent map. Thus the image of (D', π') is also the class of Δ , which shows that $\tilde{\Phi}$ is well-defined.

It remains to show that $\tilde{\Phi}$ has the stated property. For this let $Q \in D(L)$ be a point defined over some extension L of K . We can find an L -rational divisor $d = \sum_i n_i Q_i$ on D linearly equivalent to $[Q]$ and such that the support of d contains no points lying above the flex points of C or the zeros of u . The divisor $[\pi(Q)]$ on C is linearly equivalent to the L -rational divisor $\pi_* d := \sum_i n_i [\pi(Q_i)]$ (e.g. [Sil, II.3.6]). So $\Phi(\pi(Q))$ is represented by $\ell(\pi_* d)$. On the other hand, the relation (3.1) defining Δ gives,

$$\ell(\pi_* d) = \Delta \cdot \left(\frac{\partial h}{u \circ \pi} \right) (d),$$

since $\deg(d) = 1$. Now since d is L -rational, $\left(\frac{\partial h}{u \circ \pi} \right) (d) \in L^\times \partial F_L^\times$. So $\Phi(\pi(Q))$ is represented by Δ as required. \square

In what follows we will refer to a linear form $h \in F[z_1, \dots, z_{p^2}]$ as in the proof (i.e. such that $\pi^*[\mathbf{x}] = \text{div}(h)$) as a linear form defining the pull-back of a generic flex. Recall that ℓ is defined as the ratio $\ell = \tilde{\ell}/u$. If $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$ and h is a linear form defining the pull-back of a generic flex (on some model of (D, π)), then the image of (D, π) under the descent map is also represented by the $\Delta \in H^\times$ satisfying the relation

$$\tilde{\ell} \circ \pi = \Delta \partial h$$

in the coordinate ring of D .

4. Injectivity of the descent map

The fake descent map factors through the descent map as

$$\tilde{\Phi}_{fake} : \text{Cov}_0^{(p)}(C/K) \xrightarrow{\tilde{\Phi}} \frac{H^\times}{K^\times \partial F^\times} \xrightarrow{\text{pr}_1} \frac{F^\times}{K^\times F^{\times p}}.$$

So the descent map carries more information. In this section we show that this extra information is enough to eliminate any ambiguity.

PROPOSITION 4.1. *Let $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$, $\xi \in C^\perp$ and $(D, \pi) \cdot \xi$ be the twist of (D, π) by ξ . Then*

$$\tilde{\Phi}((D, \pi) \cdot \xi) = \tilde{\Phi}((D, \pi)) \cdot \tilde{\Phi}_0(\xi) \in H^\times / K^\times \partial F^\times,$$

where $\tilde{\Phi}_0 : C^\perp \simeq (\partial \bar{F}^\times)^{G_K} / K^\times \partial F^\times$ is the isomorphism given by corollary 2.11.

Recall that C^\perp acts simply transitively on $\text{Cov}_0^{(p)}(C/K)$ by twisting. The proposition shows that the image of C^\perp under $\tilde{\Phi}_0$ acts on the image of $\text{Cov}_0^{(p)}(C/K)$ under $\tilde{\Phi}$ by multiplication and that the pair $(\tilde{\Phi}, \tilde{\Phi}_0)$ respects these two actions. Since $\tilde{\Phi}_0$ is an isomorphism, we deduce the following.

COROLLARY 4.2. *Assume $\text{Cov}_0^{(p)}(C/K)$ is nonempty. Then the descent map is an affine isomorphism (i.e. isomorphism of principal homogeneous spaces). In particular $\tilde{\Phi} : \text{Cov}_0^{(p)}(C/K) \rightarrow H^\times / K^\times \partial F^\times$ is injective, and its image is a coset of $(\partial \bar{F}^\times)^{G_K} / K^\times \partial F^\times$ inside $H^\times / K^\times \partial F^\times$.*

Before giving the proof, it will be useful to put together a diagram. For any $x_0 \in X$ and $P \in E[p]$, the Weil pairing on $E[p]$ gives a map

$$\phi_{P, x_0} : X \ni x \mapsto e_p(P, x - x_0) \in \mu_p,$$

where $x - x_0$ denotes the unique $T \in E[p]$ such that $x_0 + T = x$. A different choice for x_0 gives a map which differs by a constant factor. Thus, the image of ϕ_{P, x_0} in

$$\mu_p(\bar{F}) / \mu_p = \text{Map}(X, \mu_p) / \{\text{constant maps}\}$$

depends only on P . Nondegeneracy of the Weil pairing shows that distinct choices for P lead to distinct maps. Thus we have an embedding $E[p] \hookrightarrow \mu_p(\bar{F}) / \mu_p$.

Recall that the kernel of $\partial|_{\bar{F}^\times}$ is the space of affine maps to μ_p . Since ∂_1 is just the p -th power map, the space of affine maps is also equal to the kernel of $\partial_2|_{\mu_p(\bar{F})}$. Since the constant maps are affine, ∂_2 induces a map on $\mu_p(\bar{F}) / \mu_p$. For any $P \in E[p]$ and $x_0 \in X$, the map ϕ_{P, x_0} is affine (cf. the proof of 2.4). Now, by counting dimensions for example, we see that the sequence

$$0 \rightarrow E[p] \rightarrow \mu_p(\bar{F}) / \mu_p \xrightarrow{\partial_2} \mu_p(\bar{H}_2)$$

is exact.

We also have an exact sequence

$$1 \rightarrow \text{Aff}(X, \mu_p) \rightarrow \mu_p(\bar{F}) \xrightarrow{\partial_2} \mu_p(\bar{H}_2).$$

We claim that the two are compatible in the sense that following diagram commutes. The map $\text{Aff}(X, \mu_p) \rightarrow E[p]$ is given by projecting an affine map onto its linear part

and then identifying $E[p]$ with its dual using the Weil pairing (so the vertical sequence on the left is (2.1), considered in the discussion leading up to lemma 2.4).

$$\begin{array}{ccccccc}
& & 1 & & 1 & & \\
& & \downarrow & & \downarrow & & \\
& & \mu_p & \xlongequal{\quad} & \mu_p & & \\
& & \downarrow & & \downarrow & & \\
1 & \longrightarrow & \text{Aff}(X, \mu_p) & \longrightarrow & \mu_p(\bar{F}) & \xrightarrow{\partial_2} & \partial_2(\mu_p(\bar{F})) \longrightarrow 1 \\
& & \downarrow & & \downarrow & & \parallel \\
0 & \longrightarrow & E[p] & \longrightarrow & \mu_p(\bar{F})/\mu_p & \xrightarrow{\partial_2} & \partial_2(\mu_p(\bar{F})) \longrightarrow 1 \\
& & \downarrow & & \downarrow & & \\
& & 0 & & 1 & &
\end{array}$$

Note that the rows and columns are exact.

LEMMA 4.3. *The diagram commutes.*

PROOF: We only need to show that the lower-left square commutes, the rest being obvious. Let $\phi : x \mapsto \phi(x)$ be an affine map. Choose some $x_0 \in X$. Projection onto the linear part is the map $\Lambda_\phi : P \mapsto \phi(x_0 + P)/\phi(x_0)$. Identifying $E[p]$ with its dual via the Weil pairing, Λ_ϕ is the unique $R \in E[p]$ such that, for all $P \in E[p]$, $\phi(x_0 + P)/\phi(x_0) = e_p(R, P)$. The image of this R in $\mu_p(\bar{F})/\mu_p$ is the class of the map $\phi_{R, x_0} : x \mapsto e_p(R, x - x_0)$. By the property defining R , this is equal to the map $x \mapsto \phi(x_0 + (x - x_0))/\phi(x_0) = \phi(x)/\phi(x_0)$. Modulo constant maps, this is the same as the image of ϕ in $\mu_p(\bar{F})$, so the diagram commutes. \square

REMARK: When $p = 2$ the diagram has the curious property that it is self-dual upon reflection about the obvious diagonal (e.g. $\text{Aff}(X, \mu_2)$ and $\mu_2(\bar{F})/\mu_2$ are dual as G_K -modules). For odd p , a simple dimension count shows this is no longer the case.

For the proof of the proposition, we will make use of an alternative description of the embedding $E[p] \hookrightarrow \mu_p(\bar{F})/\mu_p$.

LEMMA 4.4. *Let $D \in \text{Cov}_0^{(p)}(C/\bar{K})$ (NB: over \bar{K} , not K) and let h denote a linear form (with coefficients in \bar{F}) defining the pull-back of the generic flex point on C . For any $R \in E[p]$, the image of R under the composition $E[p] \hookrightarrow \mu_p(\bar{F})/\mu_p \hookrightarrow \bar{F}^\times/\bar{K}^\times$ is equal to the class of $\frac{h(Q+R)}{h(Q)}$, where $Q \in D$ is any point chosen so that $h(Q)$ and $h(Q+R)$ are both defined and nonzero.*

PROOF: Let $\psi : C \rightarrow E$ be the isomorphism (defined over \bar{K}) such that $p \circ \psi = \rho$ and let x_0 be the preimage of 0_E under ψ . Further, let Q_0 be any preimage of x_0 on D and $\psi_D : D \rightarrow E$ be the isomorphism defined (over \bar{K}) by $Q \mapsto (Q - Q_0)$. We have a commutative diagram,

$$\begin{array}{ccc}
E & \xleftarrow{\psi_D} & D \\
\downarrow p & & \downarrow \pi \\
E & \xleftarrow{\psi} & C
\end{array}$$

If x is a flex point, evaluating the coefficients of h at x gives a linear form h_x defining the pull-back of $[x]$ by π . Consider the function $h_x/h_{x_0} \in \kappa(\bar{D})^\times$ and its image $g_x = (\psi_D^{-1})^*(h_x/h_{x_0}) \in \kappa(\bar{E})^\times$. The divisor of h_x/h_{x_0} is $\pi^*[x] - \pi^*[x_0]$, so by commutativity $\text{div}(g_x) = p^*[(x - x_0)] - p^*[0_E]$. By definition of the Weil pairing [Sil, III.8], for any $R \in E[p]$,

$$e_p(R, x - x_0) = \frac{g_x(T + R)}{g_x(T)},$$

where $T \in E$ is any point chosen so that both numerator and denominator are defined and nonzero. Thus, we have

$$(4.1) \quad e_p(R, x - x_0) = \frac{h_x(\psi_D^{-1}(T) + R)h_{x_0}(\psi_D^{-1}(T))}{h_x(\psi_D^{-1}(T))h_{x_0}(\psi_D^{-1}(T) + R)}.$$

Considered as an element of $\bar{F}^\times = \text{Map}(X, \bar{K}^\times)$ modulo the constant maps, the right-hand side of (4.1) is represented by the map

$$\frac{h(\psi_D^{-1}(T) + R)}{h(\psi_D^{-1}(T))} = \left(x \mapsto \frac{h_x(\psi_D^{-1}(T) + R)}{h_x(\psi_D^{-1}(T))} \right).$$

On the other hand, the left-hand side of (4.1) represents the image of R in $\mu_p(\bar{F})/\mu_p$, so we are done. \square

Proof of Proposition 4.1 Let (D, π) , (D_ξ, π_ξ) and ξ be as in the proposition, and fix models for everything in \mathbb{P}^{p^2-1} . We have an isomorphism (of coverings) $\varphi : D_\xi \rightarrow D$ defined over \bar{K} , with the property that $\varphi^\sigma(Q) = \varphi(Q) + \xi_\sigma$ for all $Q \in D_\xi$ and $\sigma \in G_K$.

Choose linear forms h and h_ξ with coefficients in F defining the pull-backs of the generic flex by π and π_ξ , respectively. For some $\Delta, \Delta_\xi \in H^\times$, necessarily representing the images of (D, π) and (D_ξ, π_ξ) under $\tilde{\Phi}$, we have

$$\Delta \cdot \partial h = \tilde{\ell} \circ \pi \quad \text{and} \quad \Delta_\xi \cdot \partial h_\xi = \tilde{\ell} \circ \pi_\xi$$

in the coordinate rings of D and D_ξ , respectively. Applying φ^* to the first relation and comparing with the second gives

$$\Delta \cdot \partial(h \circ \varphi) = \Delta_\xi \cdot \partial h_\xi$$

in the coordinate ring of D_ξ . Specializing to a point Q in D_ξ not lying above any flex point of C (i.e. so that neither h_ξ nor $h \circ \varphi$ vanish at Q) we have

$$\frac{\Delta_\xi}{\Delta} = \partial \left(\frac{h(\varphi(Q))}{h_\xi(Q)} \right) \in (\partial \bar{F}^\times)^{G_K}.$$

Note that $h_\xi(Q)$ and $h(\varphi(Q))$ depend on a choice of homogeneous coordinates for Q , but that their ratio does not. This is G_K -invariant since Δ and Δ_ξ are in H^\times .

Under the isomorphism $(\partial \bar{F}^\times)^{G_K} / \partial F^\times \simeq H^1(k, \text{Aff}(X, \mu_p))$ given in lemma 2.10, $\partial \left(\frac{h(\varphi(Q))}{h_\xi(Q)} \right)$ corresponds to the class of the cocycle

$$\eta : G_K \ni \sigma \mapsto \left(\frac{h(\varphi(Q))}{h_\xi(Q)} \right)^\sigma \left(\frac{h_\xi(Q)}{h(\varphi(Q))} \right) \in \mu_p(\bar{F}) = \text{Map}(X, \mu_p),$$

which a priori takes values in $\text{Aff}(X, \mu_p) \subset \mu_p(\bar{F})$. We need to show that the image of this cocycle under the map induced by $\text{Aff}(X, \mu_p) \rightarrow E[p]$ is cohomologous to ξ . For

this we make use of the following commutative diagram

$$(4.2) \quad \begin{array}{ccccc} \text{Aff}(X, \mu_p) & \hookrightarrow & \mu_p(\bar{F}) & \hookrightarrow & \bar{F}^\times \\ \downarrow & & \downarrow & & \downarrow \\ E[p] & \hookrightarrow & \mu_p(\bar{F})/\mu_p & \hookrightarrow & \bar{F}^\times/\bar{K}^\times \end{array}$$

Since the horizontal maps are all injective, it will be enough to show that, for any $\sigma \in G_K$, the images of ξ_σ and η_σ in the lower-right corner are equal. For this we will make use of the preceding lemma.

Using the fact that h and h_ξ are defined over H and rearranging, we have

$$\eta_\sigma = \left(\frac{h(\varphi^\sigma(Q^\sigma))}{h(\varphi(Q))} \right) \left(\frac{h_\xi(Q)}{h_\xi(Q^\sigma)} \right).$$

Making use of the fact that $\varphi^\sigma(Q^\sigma) = \varphi(Q^\sigma) + \xi_\sigma$ we can rewrite this as

$$\eta_\sigma = \left(\frac{h(\varphi(Q) + \xi_\sigma + (\varphi(Q^\sigma) - \varphi(Q)))}{h(\varphi(Q))} \right) \left(\frac{h_\xi(Q^\sigma + (Q - Q^\sigma))}{h_\xi(Q^\sigma)} \right).$$

By lemma 4.4 this represents the image of

$$\xi_\sigma + (\varphi(Q^\sigma) - \varphi(Q)) - (Q^\sigma - Q) \in E[p]$$

under the embedding given by the bottom row of (4.2). But

$$(\varphi(Q^\sigma) - \varphi(Q)) - (Q^\sigma - Q) = 0_E$$

(see [Sil, X.3.5]) so the images of η_σ and ξ_σ in the lower right corner of (4.2) are equal. From this the proposition follows. \square

We can use a similar argument to prove the following useful lemma. This says that we could use ℓ to perform descent on $E = \text{Jac}(C)$. In practice, one is likely to have produced C by performing a descent on E , so this is not going to yield anything new. It does however allow us to relate the descent on C to the descent on E .

LEMMA 4.5. *The following diagram is commutative.*

$$\begin{array}{ccc} \text{Pic}_K^0(C) & \xlongequal{\quad} & E(K) \\ \downarrow \Phi & & \downarrow \delta_E \\ \frac{(\partial \bar{F}^\times)^{G_K}}{K^\times \partial F^\times} & \xrightarrow{\tilde{\Phi}_0^{-1}} & C^\perp \hookrightarrow H^1(K, E[p]) \end{array}$$

Here δ_E is the connecting homomorphism from the Kummer sequence associated to E , and the composition of the bottom row is the map identifying $(\partial \bar{F}^\times)^{G_K}/K^\times \partial F^\times$ with $C^\perp \subset H^1(K, E[p])$.

PROOF: Let $\Xi \in \text{Pic}_K^0(C)$ and choose a representative $d \in \text{Div}(C)$ whose support is disjoint from X and any zeros of u . Write d as a difference $d = d_1 - d_2$ of effective divisors and write each d_i as a sum $d_i = \sum_{j=1}^n P_{i,j}$ of $n = \deg(d_1) = \deg(d_2)$ (possibly non-distinct) points on C . Now choose any $(D, \pi) \in \text{Cov}_0^{(p)}(C/\bar{K})$ and a linear form h with coefficients in \bar{F} defining the pull-back of the generic flex. For each $P_{i,j}$ in the support of d , choose a point $Q_{i,j} \in D$ such that $\pi(Q_{i,j}) = P_{i,j}$. These choices are such that, as points on E ,

$$p(Q_{i,j} - Q_{i',j'}) = (P_{i,j} - P_{i',j'}),$$

for any i, j, i', j' . In particular, $p \sum_{j=1}^n (Q_{1,j} - Q_{2,j}) = d$. So the image of Ξ under the connecting homomorphism is given by the cocycle

$$\sigma \mapsto \left(\sum_{j=1}^n (Q_{1,j}^\sigma - Q_{2,j}^\sigma) - \sum_{j=1}^n (Q_{1,j} - Q_{2,j}) \right) \in E[p].$$

On the other hand, the image of Ξ under Φ is represented by

$$\frac{\ell(d_1)}{\ell(d_2)} = \prod_{j=1}^n \frac{\ell(P_{1,j})}{\ell(P_{2,j})}.$$

Choose homogeneous coordinates for the $P_{i,j}$ which are compatible with the action of the Galois group (i.e. so that applying σ to the coordinates of $P_{i,j}$ gives the homogeneous coordinates for $P_{i,j}^\sigma$). The class of $\Phi(\Xi)$ is then also represented by

$$\prod_j \frac{\tilde{\ell}(P_{1,j})}{\tilde{\ell}(P_{2,j})} \in H^\times,$$

where $\tilde{\ell}(P_{i,j})$ now means evaluating the linear form on the given choice of homogeneous coordinates for $P_{i,j}$.

In the coordinate ring of D (over \bar{H}) we have $\tilde{\ell} \circ \pi = \tilde{\Phi}((D, \pi)) \cdot \partial h$. We can fix forms defining the covering map π and choose homogeneous coordinates for the $Q_{i,j}$ so that the equality $\pi(Q_{i,j}) = P_{i,j}$ is also true for the coordinates chosen. Now since $\deg(d) = 0$, we have that

$$\prod_{j=1}^n \frac{\tilde{\ell}(P_{1,j})}{\tilde{\ell}(P_{2,j})} = \prod_{j=1}^n \frac{\partial h(Q_{1,j})}{\partial h(Q_{2,j})} = \partial \left(\prod_{j=1}^n \frac{h(Q_{1,j})}{h(Q_{2,j})} \right) \in \partial \bar{F}^\times,$$

whereby $h(Q_{i,j})$ means evaluating h at the given choice of coordinates.

Under the isomorphism $(\partial \bar{F}^\times)^{G_K} / K^\times \partial F^\times \simeq H^1(K, \text{Aff}(X, \mu_p)) / K^\times$, $\Phi(\Xi)$ is sent to the class of the cocycle

$$\sigma \mapsto \alpha^\sigma / \alpha,$$

where $\alpha \in \bar{F}^\times$ is any element such that $\partial \alpha$ represents $\Phi(\Xi)$. The argument above shows we may take $\alpha = \left(\prod_{j=1}^n \frac{h(Q_{1,j})}{h(Q_{2,j})} \right)$. Hence, the image of Ξ in $H^1(K, \text{Aff}(X, \mu_p)) / K^\times$ is represented by the cocycle η sending $\sigma \in G_K$ to

$$\begin{aligned} \eta_\sigma &= \left(\prod_{j=1}^n \frac{h(Q_{1,j})}{h(Q_{2,j})} \right)^\sigma \cdot \left(\prod_{j=1}^n \frac{h(Q_{2,j})}{h(Q_{1,j})} \right) \\ &= \left(\prod_{j=1}^n \frac{h^\sigma(Q_{1,j})}{h^\sigma(Q_{2,j})} \right) \cdot \left(\prod_{j=1}^n \frac{h(Q_{2,j})}{h(Q_{1,j})} \right) \\ &= \left(\prod_{j=1}^n \frac{h^\sigma(Q_{2,j} + (Q_{1,j}^\sigma - Q_{2,j}^\sigma))}{h^\sigma(Q_{2,j})} \right) \cdot \left(\prod_{j=1}^n \frac{h(Q_{1,j} + (Q_{2,j} - Q_{1,j}))}{h(Q_{1,j})} \right) \end{aligned}$$

Applying lemma 4.4 as in the proof of the proposition, to each factor appearing, we see that the image of η_σ in $H^1(K, E[p])$ is equal to the class of the cocycle

$$G_K \ni \sigma \mapsto \sum_{j=1}^n ((Q_{1,j}^\sigma - Q_{2,j}^\sigma) - ((Q_{1,j} - Q_{2,j})) \in E[p].$$

This is the same as the image under the connecting homomorphism, so the diagram commutes. \square

5. Image of the descent map

From the preceding material we know that the image of the descent map is a coset of $(\partial\bar{F}^\times)^{G_K}/K^\times\partial F^\times$ inside $H^\times/K^\times\partial F^\times$. We would like to say even more.

From here on use \mathcal{H}_K to denote the image of $\text{Cov}_0^{(p)}(C/K)$ under the descent map and \mathcal{H}_K^0 for $(\partial\bar{F}^\times)^{G_K}/K^\times\partial F^\times$. Note that 2.11 gives an isomorphism $C^\perp \simeq \mathcal{H}_K^0$. For explicit purposes one prefers to work with representatives in H^\times . We set $\tilde{\mathcal{H}}_K^0 = (\partial\bar{F}^\times)^{G_K} \subset H^\times$. To achieve the same for \mathcal{H}_K , let $P \in C(\bar{K})$ be any point that is neither a flex nor a pole of u . Define $\tilde{\mathcal{H}}_K = (\ell(P) \cdot \partial\bar{F}^\times)^{G_K}$. That this does not depend on the choice for P is shown in the proof below.

LEMMA 5.1. $\mathcal{H}_K = \tilde{\mathcal{H}}_K/K^\times\partial F^\times$

PROOF: To show that $\tilde{\mathcal{H}}_K$ does not depend on P , let $P' \in C(\bar{K})$ be any point which is neither a zero nor a pole of ℓ . Choose $(D, \pi) \in \text{Cov}_0^{(p)}(C/\bar{K})$ (NB: over \bar{K} , not K). Fixing a model for (D, π) and choosing a linear form h defining the pullback of the generic flex, we get a relation $\ell \circ \pi = \Delta(\partial h/u \circ \pi)$ in $\kappa(D \otimes_{\bar{K}} \bar{H})$. Choosing points Q, Q' lying above P and P' we get that

$$\ell(P)/\ell(P') = \frac{\partial h(Q) \cdot u(P')}{\partial h(Q') \cdot u(P)} = \left(\frac{u(P')}{u(P)} \right) \cdot \partial \left(\frac{h(Q)}{h(Q')} \right) \in \bar{K}^\times \partial\bar{F}^\times = \partial\bar{F}^\times.$$

It follows that the coset $\ell(P) \cdot \partial\bar{F}^\times$ does not depend on P . Hence neither does its G_K -invariant subset.

Clearly if $\tilde{\mathcal{H}}_K$ is nonempty, then it is a coset of $\tilde{\mathcal{H}}_K^0$. So it will suffice to show that

$$\left(\tilde{\mathcal{H}}_K \neq \emptyset \right) \implies \left(\emptyset \neq \mathcal{H}_K \subset \tilde{\mathcal{H}}_K/K^\times\partial F^\times \right).$$

In section 7 we show how to construct representatives for elements of $\text{Cov}_0^{(p)}(C/K)$ from elements of $\tilde{\mathcal{H}}_K$, so we will assume $\mathcal{H}_K \neq \emptyset$.

To show containment, let $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$. Its image in \mathcal{H}_K is defined by a relation in the coordinate ring of D of the form $\tilde{\ell} \circ \pi = \Delta \partial h$. Evaluating at any point $Q \in D$ not lying above a flex or zero of u , we see that $\Delta \in \ell(\pi(Q)) \cdot \partial\bar{F}^\times$. But we know Δ is Galois invariant, so it must lie in $(\ell(\pi(Q)) \cdot \partial\bar{F}^\times)^{G_K} = \tilde{\mathcal{H}}_K$. \square

The value of this description is that it shows that non-membership in \mathcal{H}_K is stable under base change.

COROLLARY 5.2. *Let $K \subset L$ be any extension of fields and $\Delta \in H^\times$. If $\Delta \otimes_K 1_L \in \tilde{\mathcal{H}}_L$, then $\Delta \in \tilde{\mathcal{H}}_K$. The same is true of \mathcal{H}_K .*

PROOF: At least for algebraic extensions (i.e. $L \subset \bar{K}$), this is evident from the fact that $\tilde{\mathcal{H}}_K$ is defined by taking Galois invariants. One might be able to deduce the general case from this. We give an alternate proof using geometric methods in 7.16 below. \square

This means we can work over any extension to decide whether an element of H^\times represents a class in \mathcal{H}_K . Over a number field k , this works as follows. We pick a prime

v of k and compute the images under $\tilde{\Phi}_v$ of all coverings $(D_v, \pi_v) \in \text{Cov}_0^{(p)}(C/k_v)$ with $D(k_v) \neq \emptyset$ by finding sufficiently many independent points of $C(k_v)$, up to sufficiently high precision, and evaluating ℓ on them (cf. Section III.2). By the corollary, those $\Delta \in H^\times/k^\times \partial F^\times$ that restrict into this set necessarily lie in \mathcal{H}_k . Those that do not may still lie in \mathcal{H}_k , but are certainly not in the image of the Selmer set under the descent map. So they can be ignored.

REMARK: The analog of this statement for the fake descent map fails. In the case $p = 2$ for example, in order that $\delta \in F^\times$ represent the image of some 2-covering under the fake descent map it is necessary and sufficient that $c \cdot N_{F/K}(\delta)$ be a square (see below). This means that every class in $F^\times/K^\times F^{\times 2}$ is the image of some 2-covering defined over a quadratic extension of K .

Explicit norm conditions.

The case $p = 2$. When $p = 2$, $H_2 = K$ and ∂_2 is the norm from F to K . We have seen (lemma 2.6) that

$$\begin{aligned} \tilde{\mathcal{H}}_K^0 &= \{(\delta, \varepsilon) \in F^\times \times H_2^\times : \partial_2(\delta) = \varepsilon^2\} \\ &= \{(\delta, \varepsilon) \in F^\times \times K^\times : N_{F/K}(\delta) = \varepsilon^2\}. \end{aligned}$$

Recall that the model for C is given by $u_3^2 = c \cdot f(u_1, u_2)$ with f monic and of degree 4 in u_1 . The functions, $\ell_1 = (u_1 - \theta u_2)/u_2$ and $\ell_2 = u_3/u_2^2$, used to define the descent map satisfy $c \cdot N_{F/K}(\ell_1) = \ell_2^2$. So $\tilde{\mathcal{H}}_K$ is the coset

$$\tilde{\mathcal{H}}_K = \{(\delta, \varepsilon) \in F^\times \times K^\times : c \cdot N_{F/K}(\delta) = \varepsilon^2\}.$$

The image of $\tilde{\mathcal{H}}_K$ under the projection to $F^\times/K^\times F^{\times 2}$ gives the familiar description (e.g [BS, MSS, Sta]) of the image of the fake descent map as the set

$$\mathcal{H}_K^{\text{fake}} = \{\delta \in F^\times/K^\times F^{\times 2} : c \cdot N_{F/K}(\delta) \in K^{\times 2}\}.$$

This gives very explicit descriptions of the images of the fake and true descent maps in terms of a norm condition. We would like something similar for odd p .

The case of odd p . Suppose $(\delta, \varepsilon) \in \tilde{\mathcal{H}}_K^0$. Then $(\delta, \varepsilon) = (\partial_1(\alpha), \partial_2(\alpha))$ for some $\alpha \in \bar{F}^\times$. Since $\partial_1(\alpha) = \alpha^p$ we have

$$\tilde{\mathcal{H}}_K^0 \subset \{(\delta, \varepsilon) \in F^\times \times H_2^\times : \partial_2(\delta) = \varepsilon^p\}.$$

Already for $p = 3$ this can be a proper inclusion. To see why suppose $(\delta, \varepsilon) \in F^\times \times H_2^\times$ with $\partial_2(\delta) = \varepsilon^p$. Choosing any $\alpha \in \bar{F}^\times$ which is a p -th root of δ we see that $\partial_2(\alpha) = \eta\varepsilon$ for some $\eta \in \mu_p(\bar{H}_2)$. In order that (δ, ε) be in $\tilde{\mathcal{H}}_K^0$ we must have that η lies in the image of $\partial_2|_{\mu_p(\bar{F})}$. But for $p \geq 3$, $\partial_2|_{\mu_p(\bar{F})}$ is not surjective (count dimensions).

Thus, while ε is a p -th root of $\partial_2(\delta)$, it is not necessarily the case that all p -th roots will lead to pairs in $\tilde{\mathcal{H}}_K^0$. One needs to impose extra conditions. We can achieve this for $p = 3$ by using the second norm map induced by the structure of Y_2 as a G_K -module derived from X (see section I.3). This is the map $\partial'_2 : H_2 \rightarrow F$ defined by

$$H_2 = \text{Map}_K(Y_2, \bar{K}) \ni \phi \mapsto \left(x \mapsto \prod_{y:x \in y} \phi(y)\right) \in \text{Map}_K(X, \bar{K}) = F.$$

For $p = 3$, the value of $\partial'_2(\phi)$ on a flex point x is the product of the values taken by ϕ on the lines (in Y_2) passing through x .

LEMMA 5.3. For $p = 3$,

$$\tilde{\mathcal{H}}_K^0 = \left\{ (\delta, \varepsilon) \in F^\times \times H_2^\times : \partial_2(\delta) = \varepsilon^3 \text{ and } \partial'_2(\varepsilon) = \delta \cdot \left(\frac{N_{H_2/K}(\varepsilon)}{N_{F/K}(\delta)} \right) \right\}.$$

Since $\tilde{\mathcal{H}}_K$ is a coset, we have the following description.

COROLLARY 5.4. For $p = 3$, there exist constants $\beta \in H_2^\times$ and $\beta' \in F^\times$ such that

$$\tilde{\mathcal{H}}_K = \left\{ (\delta, \varepsilon) \in F^\times \times H_2^\times : \partial_2(\delta) = \beta\varepsilon^3 \text{ and } \partial'_2(\varepsilon) = \beta'\delta \cdot \left(\frac{N_{H_2/K}(\varepsilon)}{N_{F/K}(\delta)} \right) \right\}.$$

PROOF: The proof makes use of the following two identities. For any $\phi \in \bar{F}^\times$,

$$\partial'_2\partial_2(\phi) = \phi^3 N_{F/K}(\phi) \text{ and } N_{H_2/K}\partial_2(\phi) = N_{F/K}(\phi)^4.$$

To prove these one considers intersection divisors on C . In the first identity, the value of the left-hand side at a given flex x is the product (with multiplicities) of the values of ϕ on the points lying on the lines passing through x . Every point of X other than x lies on exactly one of these lines while x lies on all four. Thus in the product, the value at x appears with multiplicity four, while the values at all other flexes occur with multiplicity one. This is evidently equal to the value of the right-hand side at x . One proves the second identity similarly.

Now suppose that $(\delta, \varepsilon) = (\partial_1(\alpha), \partial_2(\alpha)) \in \tilde{\mathcal{H}}_K^0$. Applying the identities to α we see that $\partial'_2(\varepsilon) = \delta N_{F/K}(\alpha)$ and $N_{H_2/K}(\varepsilon) = N_{F/K}(\delta) N_{F/K}(\alpha)$. Whence, $\tilde{\mathcal{H}}_K^0$ is contained in the set in the statement.

For the other inclusion suppose that (δ, ε) is in the set in the statement. Let $\alpha \in \bar{F}^\times$ be any cube root of δ . Then $\partial_2(\alpha) = \eta\varepsilon$, for some $\eta \in \mu_3(\bar{H}_2)$. It is enough to show $\eta = \partial_2(\zeta)$ for some $\zeta \in \mu_3(\bar{F})$, for then $(\delta, \varepsilon) = \partial(\alpha\zeta^{-1}) \in (\partial\bar{F}^\times)^{G_K} = \tilde{\mathcal{H}}_K^0$.

Using linear algebra over \mathbb{F}_3 one can easily check (preferably using a computer) that

$$\text{im}(\partial_2 | \mu_3(\bar{F})) = \ker(\partial'_2 - N_{H_2/K} | \mu_3(\bar{H}_2))$$

(actually, the identities above show “ \subset ” so one only needs to determine the dimension of the kernel). So we have to show $\partial'_2(\eta) = N_{H_2/K}(\eta)$.

Applying ∂'_2 to $\partial_2(\alpha) = \eta\varepsilon$ we obtain

$$\partial'_2\partial_2(\alpha) = \partial'_2(\eta)\partial'_2(\varepsilon) = \partial'_2(\eta)\delta \cdot \left(\frac{N_{H_2/K}(\varepsilon)}{N_{F/K}(\delta)} \right)$$

Using the first identity the left-hand side is equal to $\delta N_{F/K}(\alpha)$. Using the second identity the right-hand side becomes

$$\partial'_2(\eta)\delta \cdot \left(\frac{N_{H_2/K}(\varepsilon)}{N_{F/K}(\delta)} \right) = \partial'_2(\eta)\delta \cdot \left(\frac{N_{H_2/K}(\partial_2(\alpha))}{N_{H_2/K}(\eta)N_{F/K}(\alpha^3)} \right) = \delta N_{F/K}(\alpha) \cdot \left(\frac{\partial'_2(\eta)}{N_{H_2/K}(\eta)} \right).$$

Comparing this with the left hand side we conclude $\partial'_2(\eta) = N_{H_2/K}(\eta)$ as required.

To prove the corollary, we argue as follows. The divisor cut out by $\tilde{\ell}_1$ is $3[\mathbf{x}]$, while that cut out by $\tilde{\ell}_2$ is $\partial_2[\mathbf{x}]$. So in the coordinate ring of $C \otimes_K H_2$ there is a relation $\partial_2(\tilde{\ell}_1) = \beta\tilde{\ell}_2^3$, which determines $\beta \in H_2^\times$. Similarly $N_{F/K}(\tilde{\ell}_1) \cdot \partial'_2(\tilde{\ell}_2)$ and $\tilde{\ell}_1 \cdot N_{H_2/K}(\tilde{\ell}_2)$ cut out the same divisor on C . So $\beta' \in F^\times$ can be computed by taking their ratio modulo the ideal generated by the homogeneous cubic defining C . \square

For general p arguing similarly we obtain the following. These constants will be used in the algorithm presented in Chapter III.

LEMMA 5.5. *There exist constants $c \in K^\times$ and $\beta \in H_2^\times$ such that any $(\delta, \varepsilon) \in \mathcal{H}_K$ satisfies $N_{F/K}(\delta) \equiv c \pmod{K^{\times p}}$ and $\partial_2(\delta) = \beta\varepsilon^p$.*

PROOF: The existence of β is proven exactly as in the lemma above. For the first condition note that the divisor on C defined by $N_{F/k}(\tilde{\ell}_1)$ is p times the divisor $\sum_{x \in X} [x]$. The divisor $\sum_{x \in X} [x]$ is itself cut out by some form $g \in K[u_1, \dots, u_p]$ of degree p . Thus, there is some $c \in K^\times$, such that in the coordinate ring of C

$$N_{F/K}\tilde{\ell}_1 = c \cdot g^p.$$

Arguing as above, we see that all $(\delta, \varepsilon) \in \mathcal{H}_K$ satisfy the relation in the statement. \square

REMARK: In the case $p = 3$, one may take g to be the determinant of the Hessian matrix of the homogeneous cubic defining C and the choice is unique up to scalar multiple. In general there may be many choices. In any event, the class of the constant modulo p -th powers does not depend on the choice.

In principle one should be able to find more conditions which describe $\tilde{\mathcal{H}}_K$ completely. In practice, however, such a description is not actually necessary since we can apply corollary 5.2 instead.

Let \mathcal{H}_K^{fake} denote the image of $\tilde{\mathcal{H}}_K$ in $F^\times/K^\times F^{\times p}$ under the map induced by pr_1 . Then \mathcal{H}_K^{fake} is the image of $\text{Cov}_0^{(p)}(C/K)$ under the fake descent map and

$$\mathcal{H}_K^{fake} \subset \{\delta \in F^\times/K^\times F^{\times p} : \partial_2(\delta) \in \beta \cdot H_2^{\times p}\}.$$

The next lemma identifies a sufficient condition for equality.

LEMMA 5.6. *If the natural map $H^1(K, \partial_2(\mu_p(\bar{F}))) \rightarrow H^1(K, \mu_p(\bar{H}_2))$ induced by the inclusion $\partial_2(\mu_p(\bar{F})) \subset \mu_p(\bar{H}_2)$ is injective, then*

$$\mathcal{H}_K^{fake} = \{\delta \in F^\times/K^\times F^{\times p} : \partial_2(\delta) \in \beta \cdot H_2^{\times p}\}.$$

This is satisfied when $p < 5$.

PROOF: We take Galois cohomology of the exact sequence

$$1 \rightarrow \text{Aff}(X, \mu_p) \rightarrow \mu_p(\bar{F}) \xrightarrow{\partial_2} \partial_2(\mu_p(\bar{F})) \rightarrow 1.$$

This gives a commutative diagram with exact top row:

$$\begin{array}{ccccc} H^1(K, \text{Aff}(X, \mu_p)) & \longrightarrow & H^1(K, \mu_p(\bar{F})) & \longrightarrow & H^1(K, \partial_2(\mu_p(\bar{F}))) \\ \downarrow & & \downarrow & & \downarrow \\ (\partial\bar{F}^\times)^{G_K}/\partial F^\times & \longrightarrow & F^\times/F^{\times p} & \xrightarrow{\partial_2} & H_2^\times/H_2^{\times p} \end{array}$$

The left and middle vertical maps are isomorphisms given by lemma I.2.10 and HT90, respectively. Up to another application of Theorem 90, the vertical map on the right is the map in the statement. If this is injective, then the bottom row is exact. The description of \mathcal{H}_K^{fake} follows since it is a coset of the kernel of ∂_2 .

Injectivity holds trivially for $p = 2$; the map is the identity. For the general case, consider the short exact sequence

$$1 \rightarrow \partial_2(\mu_p(\bar{F})) \rightarrow \mu_p(\bar{H}_2) \rightarrow Q \rightarrow 1,$$

where Q denotes the quotient. One needs to show that the map $\mu_p(H_2) \rightarrow Q^{G_K}$ is surjective. For given p there are only finitely many possibilities (corresponding to conjugacy classes of subgroups in $\text{AGL}_2(\mathbb{F}_p)$). In the case $p = 3$, one can check (using MAGMA, for example) that in each case the map is indeed surjective. \square

For $p \geq 5$, the same computations show that injectivity can fail. In the number field case with $p < 5$, a corollary is that a local-global analog of corollary 5.2 is valid for \mathcal{H}_k^{fake} . Recall that $\text{Sel}_{fake}^{(p)}(C/k)$ was defined as the set of classes in $F^\times/k^\times F^{\times p}$ which restrict into the image of

$$\Phi_{fake,v} : \text{Pic}_{k_v}^1(C) \xrightarrow{\Phi_v} \mathcal{H}_v \xrightarrow{\text{pr}_1} F_v^\times/k_v^\times F_v^{\times p}$$

at all primes of k , whereas \mathcal{H}_k^{fake} is the image of the fake descent map

$$\tilde{\Phi}_{fake} : \text{Cov}_0^{(p)}(C/K) \rightarrow F^\times/k^\times F^\times.$$

COROLLARY 5.7. *For k a number field and p equal to 2 or 3, $\text{Sel}_{fake}^{(p)}(C/k) \subset \mathcal{H}_k^{fake}$*

PROOF: For $p < 5$, lemma 5.6 gives a description of the image of the fake descent map, both locally and globally, as the set of classes represented by elements δ such that $\partial_2(\delta)/\beta$ is a p -th power. The Grunwald-Wang theorem [CoN, IX.9.1.11] implies that an element of H_2^\times is a p -th power if it is a p -th power everywhere locally. So if δ restricts into $\mathcal{H}_{k_v}^{fake}$ for all primes v of k , it represents a class in \mathcal{H}_k^{fake} . \square

REMARK: The possibility that this might fail for larger p should probably be taken as an indication that the fake Selmer set (as we have defined it) is not the correct object to consider.

6. The main diagram

Recall the exact diagram of section 4:

$$(6.1) \quad \begin{array}{ccccccc} & & 1 & & 1 & & \\ & & \downarrow & & \downarrow & & \\ & & \mu_p & \xlongequal{\quad} & \mu_p & & \\ & & \downarrow & & \downarrow & & \\ 1 & \longrightarrow & \text{Aff}(X, \mu_p) & \longrightarrow & \mu_p(\bar{F}) & \xrightarrow{\partial_2} & \partial_2(\mu_p(\bar{F})) \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \parallel \\ 0 & \longrightarrow & E[p] & \longrightarrow & \mu_p(\bar{F})/\mu_p & \xrightarrow{\partial_2} & \partial_2(\mu_p(\bar{F})) \longrightarrow 1 \\ & & \downarrow & & \downarrow & & \\ & & 0 & & 1 & & \end{array}$$

Let \mathcal{K}_K be the finite, although somewhat unwieldy, group

$$\mathcal{K}_K \simeq \frac{H^0(K, (\partial_2(\mu_p \bar{F})))}{\partial_2 \left(H^0 \left(K, \frac{\mu_p(\bar{F})}{\mu_p} \right) \right)} \subset \frac{\mu_p(H_2)}{\partial_2 \left(H^0 \left(K, \frac{\mu_p(\bar{F})}{\mu_p} \right) \right)}.$$

Taking Galois cohomology of the diagram above we have the following.

PROPOSITION 6.1 (Main Diagram). *The following diagram is exact and commutative.*

$$\begin{array}{ccccccc} & & 1 & & 1 & & \\ & & \downarrow & & \downarrow & & \\ 1 & \longrightarrow & \mathcal{K}_K & \longrightarrow & C^\perp \simeq \mathcal{H}_K^0 & \xrightarrow{\text{pr}_1} & F^\times / K^\times F^{\times p} \xrightarrow{\partial_{2*}} H^1(K, \partial_2(\mu_p(\bar{F}))) \\ & & \parallel & & \downarrow & & \parallel \\ 1 & \longrightarrow & \mathcal{K}_K & \longrightarrow & H^1(K, E[p]) & \longrightarrow & H^1(K, \frac{\mu_p(\bar{F})}{\mu_p}) \xrightarrow{\partial_{2*}} H^1(K, \partial_2(\mu_p(\bar{F}))) \\ & & & & \downarrow \gamma & & \downarrow \\ & & & & \text{Br}(K)[p] & \longlongequal{\quad} & \text{Br}(K)[p] \end{array}$$

PROOF: The lower of the two rows is obtained directly from the long exact sequence of the bottom row of (6.1). Up to the identifications described below, the upper row is obtained by taking Galois cohomology of the upper row of (6.1) and then modding out by the images of $H^1(K, \mu_p)$. A completely formal diagram chase shows that the kernels of these two rows must be isomorphic (so that \mathcal{K}_K is the kernel in the top row as well).

The identifications are the obvious ones following from HT90, lemma 2.10 and corollary 2.11. One can check that the map labelled pr_1 is in fact induced by projection of $H \simeq F \times H_2$ onto its first factor. This is obvious from the proof of 2.11. \square

REMARK: The point of lemma 5.6 is to show that, for $p = 2, 3$, we can replace $H^1(K, \partial_2(\mu_p(\bar{F})))$ with $H_2^\times / H_2^{\times p}$ without affecting exactness. In this case, ∂_{2*} is equal to the map induced by $\partial_2 : F \rightarrow H_2$ (as it should be). In general,

$$\text{pr}_1(\mathcal{H}_K^0) \subset \{ \delta \in F^\times / K^\times F^{\times p} : \partial_2(\delta) \in H_2^{\times p} \},$$

but the inclusion may be proper.

The relevant data for descent on C is contained in a translate of the top row. For any $(\delta, \varepsilon) \in \mathcal{H}_K$ we have a commutative diagram. The lower row is an exact sequence of groups; the upper row an exact sequence of pointed sets:

$$\begin{array}{ccccccc} & & \text{Cov}_0^{(p)}(C/K) & & & & \\ & & \downarrow \tilde{\Phi} & \searrow \tilde{\Phi}_{fake} & & & \\ 1 & \longrightarrow & \mathcal{K}_K \cdot (\delta, \varepsilon) & \longrightarrow & \mathcal{H}_K & \xrightarrow{\text{pr}_1} & \ker(\partial_{2*}) \cdot \delta \longrightarrow 1 \\ & & \cdot(\delta, \varepsilon) \uparrow & & \cdot(\delta, \varepsilon) \uparrow & & \cdot \delta \uparrow \\ 1 & \longrightarrow & \mathcal{K}_K & \longrightarrow & \mathcal{H}_K^0 & \xrightarrow{\text{pr}_1} & \ker(\partial_{2*}) \longrightarrow 1 \end{array}$$

Suppose the image of $(D, \pi) \in \text{Cov}_0^{(p)}(C/K)$ under the descent map is (δ, ε) . Then the image under the fake descent map is δ . There are $\#\mathcal{K}_K$ isomorphism classes of p -coverings with trivial obstruction that map to the class of δ under the fake descent map. Their images under $\tilde{\Phi}$ are represented by the pairs $\partial(\eta) \cdot (\delta, \varepsilon) = (\delta, \partial_2(\eta)\varepsilon)$, as $\partial_2(\eta) \in (\partial_2(\mu_p(\bar{F})))^{G_K}$ ranges over a set of representatives for \mathcal{K}_K .

When $p = 2$, $H_2 \simeq K$ and $\partial_2 : F \rightarrow K$ is the norm. A map $\phi \in \mu_2(\bar{F}) = \text{Map}(X, \mu_2)$ gives a Galois invariant class in $\mu_2(\bar{F})/\mu_2$ if and only if ϕ takes a constant value on each G_K orbit in X . The norm of such a map can be non-trivial only if there is some orbit of odd order. This can happen only if there is a K -rational flex point. This shows that \mathcal{K}_K is either of order 2 or trivial, correspondingly as $X(K)$ is or is not empty. We recover the fact that the ‘fake descent map’ is two to one unless C is trivial as a 2-covering of its Jacobian, in which case the map is injective. This is a special case of the following lemma.

LEMMA 6.2. *If there exists a K -rational flex point, then $\mathcal{K}_K = 0$.*

PROOF: Using the K -rational flex as a base point for an Abel-Jacobi map to the Jacobian gives a K -isomorphism $C \simeq E$ identifying X and $E[p]$ as G_K -sets. The functions used to define the fake descent map are the same as those used to perform p -descent on E (as described in section I.4). So injectivity follows from the corresponding statement from p -descent on elliptic curves. \square

REMARK: Unlike the case for $p = 2$, the converse of this statement does not hold when p is odd. There are examples where G_K acts transitively on X and $\mathcal{K}_K = 0$.

In general, \mathcal{K}_K can be computed if one knows the Galois action on the flex points. As previously mentioned, the action factors through the affine general linear group, $\text{AGL}_2(\mathbb{F}_p)$ (and \mathcal{K}_K depends only on the conjugacy class of the image). For fixed p there are only finitely many possibilities. The following table gives some indication of the situation for small p . The entries give the number of subgroups of $\text{AGL}_2(\mathbb{F}_p)$, up to conjugacy, for which $\dim_{\mathbb{F}_p}(\mathcal{K}_K)$ is larger than the indicated value. For example, in the case $p = 3$, there are 46 subgroups up to conjugacy and $\#\mathcal{K}_K \in \{1, 3, 9, 27\}$. For $46 - 13 = 33$ of the possible Galois actions (including the generic situation) \mathcal{K}_K is trivial.

p	$\dim_{\mathbb{F}_p} \mathcal{K}_K \geq 0$	≥ 1	≥ 2	≥ 3	≥ 4
2	11	7	0	0	0
3	46	13	4	2	0
5	132	28	7	3	0
7	236	37	4	2	0

7. Inverse of the descent map

The main result of this section is the explicit construction of an inverse to the descent map. Recall that $\tilde{\mathcal{H}}_K \subset H^\times$ is the subset of elements which represent classes in the image of $\tilde{\Phi} : \text{Cov}_0^{(p)}(C/K) \rightarrow H^\times/K^\times \partial F^\times$.

THEOREM 7.1. *Given $\Delta \in \tilde{\mathcal{H}}_K$, we can explicitly compute a set of $p^2(p^2 - 3)/2$ linearly independent quadrics over K which define a genus one normal curve $D_\Delta \subset \mathbb{P}^{p^2-1}$ of degree p^2 and a set of homogeneous polynomials defining a map $\pi_\Delta : D_\Delta \rightarrow C$*

making D_Δ into a p -covering of C . Moreover, the image of the class of (D_Δ, π_Δ) in $\text{Cov}_0^{(p)}(C/K)$ under the descent map is equal to the class of Δ in \mathcal{H}_K .

We are going to give two proofs of this theorem (at least for odd p). The first is strongly influenced by [CFOSS-II, Section 3]. In that paper, the problem of obtaining models of n -coverings of elliptic curves with trivial obstruction as genus one normal curves of degree n is considered. Their first step (in the ‘Segre embedding method’) is to embed the curve as a genus one normal curve of degree n^2 . They then show that, after projection to a suitable hyperplane, the embedding factors through the Segre embedding $\mathbb{P}^{n-1} \times (\mathbb{P}^{n-1})^\vee \rightarrow \mathbb{P}^{n^2-1}$. Making this factorization explicit requires an explicit trivialization of the obstruction algebra associated to the n -covering.

In our situation, things are actually somewhat simpler. We start with a p -covering of C . The analog of the first step of the Segre embedding method above yields a model as genus one normal curve of degree p^2 . This already gives us what we are after. It is a feature of second p -descents that no trivialization of the obstruction algebra is necessary. One will note that this is also the case for $p = 2$.

Our second proof is (on its own) actually incomplete; it seems to leave open the possibility that the scheme D_Δ that we construct from a given $\Delta \in \tilde{\mathcal{H}}_K$ splits into p components each of which is a p -isogeny covering² of C . The primary reason for including it here is that it yields a proof of 5.2. Namely, one wants to know what the construction yields if one starts with $\Delta \in H^\times \setminus \tilde{\mathcal{H}}_K$. Not surprisingly, one gets a 0-dimensional scheme.

Another advantage to the second proof is that it gives a geometric interpretation of the relation between the fake and genuine descent maps. In order to show that D_Δ is a p -covering, we first work only with δ , where $\Delta = (\delta, \varepsilon)$. To it we associate a model in \mathbb{P}^{p^2-1} for a (neither smooth nor irreducible) curve D_δ of degree p^{p^2-1} with a map to C . In a sense, these D_δ are the geometric objects witnessed by the fake descent map. The group $H^1(K, \mu_p(\bar{F})/\mu_p)$ appearing in the main diagram can be interpreted as parameterizing the twists of such objects, giving some explanation for its appearance there.

An important theoretical consequence of this theorem that should not be overlooked is that $\text{Cov}_0^{(p)}(C/K)$ is non-empty whenever $\tilde{\mathcal{H}}_K$ is. This completes the proof of lemma 5.1.

The case $p = 2$. We briefly recall the method described in [MSS, Sta, BS] for constructing the coverings in the case $p = 2$. The only, very minor, difference here is that we work with \mathcal{H}_K rather than its image in $F^\times/K^\times F^{\times 2}$.

Recall that we have a model for C (in $(1, 1, 2)$ -weighted projective space) given by an equation $u_3^2 = c \cdot f(u_1, u_2)$, where $f(u_1, u_2)$ is a binary quartic, monic in u_1 , and $c \in K^\times$. The classes in \mathcal{H}_K are represented by the elements in

$$\tilde{\mathcal{H}}_K = \{(\delta, \varepsilon) \in F^\times \times K^\times : c \cdot N(\delta) = \varepsilon^2\},$$

where $N = N_{F/K}$ is the norm. To construct the covering corresponding to (δ, ε) , one considers the equation

$$\tilde{\ell} = (\delta, \varepsilon) \cdot a\partial(z),$$

²By p -isogeny covering we mean a étale covering that is geometrically Galois with cyclic Galois group isomorphic to the kernel of some degree p isogeny on the Jacobian.

where we think of $a \in K^\times$ and $z \in F \setminus \{0\}$ as unknowns. Note that this really corresponds to the two equations

$$(7.1) \quad u_1 - \theta u_2 = \delta a z^2 \text{ and } u_3 = \varepsilon a^2 N(z),$$

with coefficients in F and K , respectively.

REMARK: Typically one works with elements of

$$\{\delta \in F^\times : c \cdot N(\delta) \in K^{\times 2}\},$$

which give only the first equation. To get the second equation one needs to choose a square root of $c \cdot N(u_1 - \theta u_2) = c \cdot N(\delta)N(az^2)$. Note that the fiber over δ in $\tilde{\mathcal{H}}_K$ is $\{(\delta, \pm\varepsilon)\}$, so the only role played by ε is to fix this choice.

In terms of the basis $\{1, \theta, \theta^2, \theta^3\}$ we can write an arbitrary element of F as $z = \sum_{i=1}^4 z_i \theta^{i-1}$. This also allows us to identify $(F \setminus \{0\})/K^\times$ with the K -points of \mathbb{P}^3 . Under this correspondence $0 \neq z = \sum z_i \theta^{i-1} \in \bar{F}$ corresponds to the point $(z_1 : \dots : z_4) \in \mathbb{P}^3$.

Writing the first equation in (7.1) above in terms of the basis for F over K and collecting powers of θ gives four equations with coefficients in K :

$$\begin{aligned} u_1 &= a \cdot Q_1(z_1, \dots, z_4), \\ u_2 &= a \cdot Q_2(z_1, \dots, z_4), \\ 0 &= a \cdot Q_3(z_1, \dots, z_4), \\ 0 &= a \cdot Q_4(z_1, \dots, z_4), \end{aligned}$$

where the Q_i are homogeneous of degree 2. The last two equations define a quadric intersection $D_\delta \subset \mathbb{P}^3$. The only role played by a is to deal with the fact that we want to work with projective coordinates. In particular any pairs (δ, ε) and (δ', ε') which are congruent modulo K^\times will yield the same curve. The equations

$$\begin{aligned} u_1 &= Q_1(z_1, \dots, z_4), \\ u_2 &= Q_2(z_1, \dots, z_4), \\ u_3 &= \varepsilon N(z_1, \dots, z_4). \end{aligned}$$

define a morphism $\pi_{(\delta, \varepsilon)} : \mathbb{P}^3 \rightarrow \mathbb{P}^2(1, 1, 2)$. Using the fact that $c \cdot N(\delta) = \varepsilon^2$, one can check this gives a morphism $\pi_{(\delta, \varepsilon)} : D_\delta \rightarrow C$. One can show that this is in fact a 2-covering of C in various ways. But note already that if it is a 2-covering of C , then it represents a class in $\text{Cov}_0^{(2)}(C/K)$ (we have exhibited a model) and by construction its image under the descent map will be given by (δ, ε) . We outline a method that is similar in some respects to one of the approaches taken below for odd p .

First one shows that $\pi_{(\delta, \varepsilon)}$ is an unramified morphism of degree 4, and in particular finite, so D_δ is of dimension 1. For this it suffices to work geometrically. Over \bar{K} all flex points are defined, so in an appropriate choice of basis for \bar{F} over \bar{K} , the equation $u_1 - \theta u_2 = \delta a z^2$ can be written as four equations $u_1 - \theta_i u_2 = a \delta_i z_i^2$ where for $i \in \{1, \dots, 4\}$, $u_1 - \theta_i u_2 = 0$ defines the tangent line to C at the flex point $(\theta_i : 0 : 1)$.

Let $P = (P_1 : P_2 : P_3)$ be any choice of coordinates defining a point in $C(\bar{K})$. Then $P_1 - \theta_i P_2 = 0$ if and only if P represents the flex point $(\theta_i : 1 : 0)$. The system of equations

$$P_1 - \theta_i P_2 = \delta_i z_i^2; i = 1, \dots, 4$$

has 2^4 solutions (resp. 2^3 solutions) in $\bar{F} \simeq \prod_{i=1, \dots, 4} \bar{K}$ when P is not a flex point (resp. is a flex point). The additional condition $P_3 = \varepsilon N_{F/K}(z_1, \dots, z_4)$ is satisfied by exactly half (resp. all) of these. In projective coordinates, this gives $2^2 = 4$ preimages of P on D_δ . We conclude that $\pi_{(\delta, \varepsilon)}$ is unramified of degree 4. By Riemann-Hurwitz, D_δ is a smooth genus one curve.

One must also show that the covering is Galois with group $E[2]$. For this it also suffices to work geometrically. Namely, it is enough to show that over \bar{K} the corresponding extension of function fields is Galois with group isomorphic to $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ (see, for example, [Sta, 5.1.2]). For this, choose any $i_0 \in \{1, \dots, 4\}$. Then, for $i \neq i_0$,

$$T_i := \frac{u_1 - \theta_i u_2}{u_1 - \theta_{i_0} u_2} \in \kappa(\bar{C})^\times$$

is a non-constant rational function in $\kappa(\bar{C})^\times$. One can see from the defining equations that the extension of function fields corresponding to the coverings is given by

$$\kappa(\bar{C}) \hookrightarrow \kappa(\bar{C})(\sqrt{T_i} : i \neq i_0) \simeq \kappa(\bar{D}_\delta),$$

i.e. by adjoining square roots of the T_i to $\kappa(\bar{C})^\times$. However, the T_i are not independent. The norm condition (or comparing degrees) shows that the product of all three is a square in $\kappa(\bar{C})^\times$. So the extension is biquadratic, hence Galois with the appropriate Galois group.

REMARK: An alternative to this last argument is to use that the action of $\mu_2(\bar{F})/\mu_2$ on $(\bar{F} \setminus \{0\})/\bar{K}^\times$ induces an action on \mathbb{P}^3 by linear automorphisms. Under this action, $E[2] \simeq \ker(N_{F/K} | \mu_2(\bar{F})/\mu_2)$ leaves the space of solutions to $\tilde{\ell} = (\delta, \varepsilon) \cdot \partial(z)$ invariant. So this gives the action of $E[2]$ on D_δ by automorphisms compatible with $\pi_{(\delta, \varepsilon)}$. One can check directly that this action is simply transitive on each fiber.

REMARK: If one starts with $(\delta, \varepsilon) \in H^\times \setminus \tilde{\mathcal{H}}_K$, then the construction produces a quadric intersection D_δ and a map $D_\delta \rightarrow \mathbb{P}^2(1, 1, 2)$, but this will not be a map to C (cf. Corollary 5.2).

The case $p > 2$. Let $\Delta \in H^\times$. Note that we are not yet assuming $\Delta \in \tilde{\mathcal{H}}_K$. We can associate to Δ a C -scheme D_Δ defined over K as follows.

Fix a basis $\{e_1, \dots, e_{p^2}\}$ for $F = \text{Map}_K(X, \bar{K})$ over K . We can then write an arbitrary element $z \in F$ as $z = \sum_i z_i e_i$. The choice of basis gives an identification of $(F \setminus \{0\})/K^\times$ with the K -points of \mathbb{P}^{p^2-1} , $0 \neq z = \sum z_i e_i$ corresponding to the point $(z_1 : \dots : z_{p^2}) \in \mathbb{P}^{p^2-1}$.

We start with the equation

$$\tilde{\ell} = a\Delta\partial(z)$$

where $a \in K^\times$ and $z \in F \setminus \{0\}$ are treated as unknown. The map $\partial : F \rightarrow H$ can be written as a homogeneous polynomial of degree p in the z_i and so our equation corresponds to an equation

$$\tilde{\ell}(u_1, \dots, u_p) = a\Delta\partial(z_1, \dots, z_p),$$

where $\tilde{\ell}$ and ∂ are homogeneous polynomials of degrees 1 and p , respectively, both with coefficients in H . Writing this out in a basis for H over K (extending the basis chosen

above) and equating coefficients on the basis vectors gives a system of $m := [H : K]$ equations, with coefficients in K , of the form:

$$(7.2) \quad \text{linear in } u_1, \dots, u_p = \text{degree } p \text{ in } z_1, \dots, z_{p^2} .$$

We claim that the matrix defined by the coefficients of the m linear forms in the left-hand side of (7.2) has full rank (i.e. rank equal to p). For this one uses a geometric argument; over \bar{K} , ℓ splits as a tuple of linear forms defining the hyperplanes in Y and these span the space of all linear forms in $\bar{K}[u_1, \dots, u_p]$.

Eliminating u_1, \dots, u_p gives a system of equations:

$$\left\{ \begin{array}{l} u_i = a \cdot \pi_i(z_1, \dots, z_{p^2}), \quad \text{for } i = 1, \dots, p \\ 0 = a \cdot P_i(z_1, \dots, z_{p^2}), \quad \text{for } i = 1, \dots, m - p \end{array} \right\}$$

where π_i and P_i are homogeneous of degree p with coefficients in K . Let

$$\{Q_j(u_1, \dots, u_p) : j = 1, \dots, N\}$$

be the homogeneous polynomials defining the model for C as a genus one normal curve of degree p in \mathbb{P}^{p-1} . Recall that in the case $p = 3$, $N = 1$ and Q_1 is of degree 3. For larger p , $N = \frac{p(p-3)}{2}$ and the Q_j are of degree 2. We define $D_\Delta \subset \mathbb{P}^{p^2-1}(z_1 : \dots : z_{p^2})$ as the (reduced) K -subscheme defined by the vanishing of the polynomials in the set

$$\{P_i : 0 < i \leq m - p\} \cup \{Q_j(\pi_1, \dots, \pi_p) : 0 < j \leq N\}.$$

The second set is included to ensure that the rational map

$$\mathbb{P}^{p^2-1}(z_1 : \dots : z_{p^2}) \rightarrow \mathbb{P}^{p-1}(u_1 : \dots : u_p)$$

defined by $u_i = \pi_i(z_1, \dots, z_{p^2})$ restricts to a morphism $\pi_\Delta : D_\Delta \rightarrow C$. Note also that if $\Delta \equiv \Delta' \pmod{K^\times}$, then $(D_\Delta, \pi_\Delta) = (D_{\Delta'}, \pi_{\Delta'})$. In other words, (D_Δ, π_Δ) only depends on the class of Δ in H^\times/K^\times .

REMARK: In the case $p = 3$, $m = [H : K] = 21$, so $D_\Delta \subset \mathbb{P}^8$ is defined by 18 cubics and one form of degree 9. For $p > 3$, we have $m = p^2(p^2 + 1)/2$, so the model is given by $\frac{p^2(p^2+1)}{2} - p$ forms of degree p and $N = \frac{p(p-3)}{2}$ forms of degree $2p$. We will see below how to obtain a set of $\frac{p^2(p^2-3)}{2}$ quadrics generating the homogeneous ideal.

One obvious, but important, property of the construction is given in the following lemma. This says that if (D_Δ, π_Δ) is a p -covering of C , then its image under the descent map is necessarily given by Δ .

LEMMA 7.2. *If there is some $R \in C(K)$ such that $\ell(R) \in \Delta \cdot K^\times \partial F^\times$, then there exists some $Q \in D_\Delta(K)$ such that $\pi_\Delta(Q) = R$.*

PROOF: Suppose $R \in C(K)$ is such that $\ell(R) = a\Delta\partial(Q)$ with $a \in K^\times$ and $Q \in F^\times$. Choose homogeneous coordinates $(R_1 : \dots : R_p)$ for R and write $Q = \sum_i e_i Q_i$ with $Q_i \in K$. Recall that $\ell(R) = \frac{\tilde{\ell}(R_1, \dots, R_p)}{u(R_1, \dots, R_p)}$, where u is a linear form not vanishing at (R_1, \dots, R_p) . Then $\tilde{\ell}(R_1, \dots, R_p) = au(R_1, \dots, R_p)\Delta\partial(Q_1, \dots, Q_{p^2})$. Eliminating as in the construction we see that

$$\begin{aligned} R_i &= au(R_1, \dots, R_p) \cdot \pi_i(Q_1, \dots, Q_{p^2}), \quad \text{for } i = 1, \dots, p, \\ 0 &= au(R_1, \dots, R_p) \cdot P_j(Q_1, \dots, Q_{p^2}), \quad \text{for } j = 1, \dots, m - p. \end{aligned}$$

Note that $au(R_1, \dots, R_p) \in K^\times$. Since $R \in C$, the equations above say that the point $(Q_1 : \dots : Q_{p^2})$ lies in $D(K)$ and is mapped via π_Δ to R . \square

The association $H^\times \ni \Delta \mapsto (D_\Delta, \pi_\Delta)$ depends on the choice of basis for F over K . We assume all (D_Δ, π_Δ) are constructed using the same basis (and so live in the same copy of \mathbb{P}^{p^2-1}). A different choice of basis leads to objects which differ only by a linear automorphism of the ambient space. It is to be understood that this automorphism is applied to each (D_Δ, π_Δ) if we change the basis.

When working geometrically, it will be convenient to use the basis of \bar{F} over \bar{K} given by the characteristic functions. These are the maps $e_x \in \bar{F} = \text{Map}(X, \bar{K})$ (indexed by $x \in X$) taking the value 1 at x and the value 0 at all $x' \neq x$. In terms of this basis, $0 \neq z \in \bar{F} \setminus \{0\}$ corresponds to the point $z = (z_x) \in \mathbb{P}^{p^2-1}$ with z_x -coordinate given by the value of z at x . We can extend to a basis for \bar{H} over \bar{K} by taking the characteristic functions on Y and identifying $x \in X$ with the hyperplane in Y cutting out the divisor $p[x]$ on C . Then $\partial(z)$ splits as the tuple of polynomials, (indexed by $y \in Y$)

$$\partial(z) = \left(\prod_{x \in y} z_x \right)_{y \in Y},$$

where as usual the product is to be taken with appropriate multiplicities.

LEMMA 7.3. *Given $\Delta \in H^\times$ we can explicitly compute a set of $p^2(p^2 - 3)/2$ linearly independent quadrics over K which lie in the homogeneous ideal of $D_\Delta \subset \mathbb{P}^{p^2-1}$.*

REMARK: We are not (yet) claiming that these quadrics define D_Δ ; we have also not assumed that $\Delta \in \tilde{\mathcal{H}}_K$.

PROOF: Under the splitting $H \simeq F \times H_2$, write $\Delta = (\Delta_1, \Delta_2)$. The equation $\tilde{\ell} = a\Delta\partial(z)$ corresponds to the two equations

$$(7.3) \quad \tilde{\ell}_1 = a\Delta_1 z^p \text{ and } \tilde{\ell}_2 = a\Delta_2 \partial_2(z).$$

First consider the case $p \geq 5$. Recall that, as a G_K -set $Y_2 \simeq X \times \frac{E[p] \setminus \{0_E\}}{\{\pm 1\}}$ and that F may be viewed as a subalgebra of H_2 . The hyperplanes in Y_2 cut out divisors on C of the form $(p-2)[x] + [x+P] + [x-P]$. So there is a quadratic form \tilde{N} such that $\partial_2(z) = z^{p-2}\tilde{N}(z)$. We can obtain a homogeneous equation in H_2 by taking the ratio of the two equations in (7.3) and multiplying through by z^2 . We get

$$(7.4) \quad \frac{\tilde{\ell}_2}{\tilde{\ell}_1} \cdot z^2 = \left(\frac{\Delta_2}{\Delta_1} \right) \cdot \tilde{N}(z).$$

To achieve the same when $p = 3$, recall that F corresponds to the G_K -set X consisting of the 9 flex points while H_2 corresponds to the G_K -set Y_2 consisting of the 12 lines in \mathbb{P}^2 passing through three distinct flex points. We can no longer view F as a subalgebra of H_2 . Instead we work with the étale algebra $M = \text{Map}_K(Z, \bar{K})$ associated to the G_K -set Z consisting of all pairs $(x, y) \in X \times Y_2$ such that $x \in y$. Each flex is contained in four lines, while each line passes through three flexes so

$$[M : F] = 4 \text{ and } [M : H_2] = 3.$$

Recall (see section I.3) that the ‘induced norm’

$$\partial = \partial_1 \times \partial_2 : F \rightarrow F \times H_2$$

is given by $\partial_1(z) = z^3$ and $\partial_2(z) = N_{M/H_2}(z)$. So, identifying z with its image in M , we can write $\partial_2(z) = z\tilde{N}(z)$ for some quadratic form \tilde{N} . Over M we can write our

equations as

$$\tilde{\ell}_1 = a\Delta_1 z^3 \text{ and } \tilde{\ell}_2 = a\Delta_2 z\tilde{N}(z).$$

Again we can obtain a homogeneous equation in M by taking the ratio. We get an equation

$$\frac{\tilde{\ell}_2}{\tilde{\ell}_1} \cdot z^2 = \left(\frac{\Delta_2}{\Delta_1} \right) \cdot \tilde{N}(z).$$

Formally this is exactly what was obtained for $p \geq 5$. Note also that $[M : K] = 3^2(3^2 - 1)/2$, while for larger p we have $[H_2 : K] = p^2(p^2 - 1)/2$. So in either case, writing the equation out in terms of the basis over K gives $p^2(p^2 - 1)/2$ quadrics some of whose coefficients are rational functions on C . These can be eliminated using linear algebra over K to obtain a set of quadrics with coefficients in K which vanish on D_Δ .

We want to count the number of independent quadrics left after eliminating. For this we may work geometrically. For $p \geq 5$ we can index the elements of Y by pairs $(x, P) \in X \times \frac{E[p]}{\{\pm 1\}}$. The linear form $\tilde{\ell}$ splits over \bar{K} as $\tilde{\ell} = (\tilde{\ell}_{(x,P)})$, where $\tilde{\ell}_{(x,P)}$ is a linear form with coefficients in \bar{K} defining the hyperplane whose intersection with C is given by the divisor $(p-2)[x] + [x+P] + [x-P]$. Note that $P = 0 = 0_E$ is allowed. For $p = 3$ we can do the same, but with the caveat that the indexing is no longer unique. Namely, each line $y \in Y_2$ corresponds to three pairs $(x, P) \in X \times \frac{E[p]}{\{\pm 1\}}$ (we get one pair for each $x \in y$). In any case, we can still use the index (x,P) to denote the factor of \bar{H} corresponding to the line in \mathbb{P}^2 whose intersection with C is given by the divisor $[x] + [x+P] + [x-P]$.

The notation is such that for distinct $(x, P) \in X \times \frac{E[p] \setminus \{0\}}{\{\pm 1\}}$, we have distinct rational functions

$$G_{(x,P)} := \frac{\tilde{\ell}_{(x,P)}}{\tilde{\ell}_{(x,0)}} \in \kappa(\bar{C})^\times$$

with divisors $\text{div}(G_{(x,P)}) = [x+P] + [x-P] - 2[x]$. Over \bar{K} , we can work with the basis of \bar{F} given by the characteristic functions and use (z_x) for coordinates on \mathbb{P}^{p^2-1} . In terms of these and the $G_{(x,P)}$ the homogeneous equation (7.4) corresponds to a system of equations

$$(7.5) \quad G_{(x,P)} \cdot z_x^2 = \tilde{\Delta}_{(x,P)} \cdot z_{x+P} z_{x-P},$$

parameterized by $(x, P) \in \frac{E[p] \setminus \{0\}}{\{\pm 1\}}$, where for simplicity we have denoted $\Delta_{(x,P)}/\Delta_{(x,0)}$ by $\tilde{\Delta}_{(x,P)}$.

For fixed x , the $(p^2 + 1)/2$ linear forms $\tilde{\ell}_{(x,P)}$, with $P \in E[p]/\{\pm 1\}$, all define hyperplanes meeting C in x with multiplicity at least $p-2$. This gives $p-2$ nontrivial relations among them. The matrix given by the coefficients of the $\tilde{\ell}_{(x,P)}(u_1, \dots, u_p)$ has rank $\leq p - (p-2) = 2$. On the other hand, the rank must be greater than one since these do not all define the same hyperplane. This introduces a dependence among the $G_{(x,P)}$. Alternatively one can argue that the functions $G_{(x,P)}$ are all in the Riemann-Roch space $\mathcal{L}(2[x])$ which has dimension 2.

In any event, if we fix $P_0 \in E[p] \setminus \{0\}$, then for any $P \in \frac{E[p] \setminus \{0\}}{\{\pm 1\}}$, we can find $a_P, b_P \in \bar{K}$ such that

$$G_{(x,P)} = a_P G_{(x,P_0)} + b_P.$$

Using this to eliminate the $G_{(x,P)}$ from (7.5) we obtain a set of quadrics

$$a_P \tilde{\Delta}_{(x,P_0)} \cdot z_{x+P_0} z_{x-P_0} + b_P \cdot z_x^2 = \tilde{\Delta}_{(x,P)} \cdot z_{x+P} z_{x-P},$$

parameterized by $P \in \frac{E[p] \setminus \{0, \pm P_0\}}{\{\pm 1\}}$ and with coefficients in \bar{K} . Since $\tilde{\Delta}_{(x,P)} \neq 0$, these are necessarily independent. Note also that the monomials appearing in these quadrics are all of the form $z_{x+Q}z_{x-Q}$ for some $Q \in E[p]/\{\pm 1\}$. A different choice for x leads to quadrics involving a disjoint set of monomials. So, in total this gives a set of $\#X \cdot \# \left(\frac{E[p] \setminus \{0, \pm P_0\}}{\{\pm 1\}} \right) = p^2(p^2 - 3)/2$ independent quadrics as required. \square

There is an obvious action of \bar{F}^\times on $(\bar{F} \setminus \{0\})/\bar{K}^\times$ by multiplication. The choice of basis gives an identification of the latter with the \bar{K} -points of \mathbb{P}^{p^2-1} and hence a representation

$$\bar{F}^\times \ni \alpha \mapsto \varphi_\alpha \in \text{PGL}_{p^2} = \text{Aut}(\mathbb{P}^{p^2-1}).$$

Working with the basis of \bar{F} given by the characteristic functions, the representation takes the particularly simple form $\alpha = (\alpha_x) \mapsto \text{Diagonal}(\alpha_x)$; this is just coordinate-wise multiplication. Assuming we are working with a basis for F over K , we see that for any extension of fields K'/K ,

$$\varphi_\alpha \in \text{PGL}_{p^2}(K') \Leftrightarrow \alpha \in (F \otimes_K K')^\times.$$

LEMMA 7.4. *For any $\Delta \in \bar{H}^\times$ and $\alpha \in \bar{F}^\times$, the action of α on \mathbb{P}^{p^2-1} induces an isomorphism (of C -schemes) $\varphi_\alpha : D_{\partial(\alpha)\Delta} \rightarrow D_\Delta$.*

COROLLARY 7.5. *Let $\Delta \in H^\times$ and (D_Δ, π_Δ) be the corresponding C -scheme. The coset $\Delta \mathcal{H}_K^0 \subset H^\times/K^\times \partial F^\times$ parameterizes a set of twists of (D_Δ, π_Δ) as a C -scheme defined over K up to K -isomorphism.*

PROOF: To prove the lemma, use that $D_{\partial(\alpha)\Delta}$ is defined by the equation $\tilde{\ell} = \partial(\alpha)\Delta\partial(z)$. If $Q \in D_{\partial(\alpha)\Delta}$ is any point mapping to, say $P \in C$, then the point $\alpha Q \in \mathbb{P}^{p^2-1}$ evidently satisfies

$$\Delta\partial(\alpha Q) \equiv \Delta\partial(\alpha)\partial(Q) \equiv \tilde{\ell}(P) \pmod{K^\times}.$$

The equivalence here is meant for any choices of coordinates for P and Q . This means αQ is a point of D_Δ lying above P . This proves the lemma.

The lemma implies that if $\Delta \in H^\times$, then the C -schemes corresponding to the elements of $\Delta \tilde{\mathcal{H}}_K^0 = \Delta(\partial \bar{F}^\times)^{G_K}$ are all twists of (D_Δ, π_Δ) . The isomorphism $\varphi_\alpha : D_{\partial(\alpha)\Delta} \rightarrow D_\Delta$ is defined over K if and only if $\alpha \in F^\times$ in which case $\partial(\alpha)\Delta$ and Δ differ by an element of $K^\times \partial F^\times$. So $\Delta \mathcal{H}_K^0$ parameterizes the corresponding twists in $\Delta \tilde{\mathcal{H}}_K^0$ up to K -isomorphism. \square

By definition, any twist of a p -covering is a p -covering, so we can reduce to the geometric situation. To prove the theorem, it is enough to show that there is some $\Delta \in \tilde{\mathcal{H}}_{\bar{K}}$ such that (D_Δ, π_Δ) is a p -covering of C defined over \bar{K} . The proof of the following lemma also shows that for $\Delta \in \tilde{\mathcal{H}}_K$, the $p^2(p^2 - 3)/2$ quadrics obtained in lemma 7.3 generate the homogenous ideal of D_Δ .

LEMMA 7.6. *There exists some $\Delta \in \tilde{\mathcal{H}}_{\bar{K}}$ such that (D_Δ, π_Δ) is a p -covering of C .*

PROOF: For this we may work over \bar{K} , using the basis given by the characteristic functions and z_x for coordinates on \mathbb{P}^{p^2-1} . Choosing any flex point $x_0 \in X$ as origin, we may consider (C, x_0) as an elliptic curve over \bar{K} . Denote the multiplication by p map on (C, x_0) by $\pi : C \rightarrow C$. This is a p -covering of C . We are going to find some $\Delta \in \bar{H}^\times$ representing the image of (C, π) under the descent map and then show that the scheme D_Δ produced by the construction above is equal to the image of (C, π) under a certain embedding into \mathbb{P}^{p^2-1} as a genus one normal curve of degree p^2 .

To compute the image of (C, π) under the descent map we use the definition. Namely, we embed C in \mathbb{P}^{p^2-1} in such a way that the pull-back of any flex point is a hyperplane section. This amounts to finding a basis for the Riemann-Roch space of the divisor $\pi^*[x_0]$. For each $x \in X$, we can find a rational function $G_x \in \kappa(\bar{C})^\times$ with divisor $\text{div}(G_x) = \pi^*[x] - \pi^*[x_0]$. For existence of these functions, note that π is multiplication by p on the elliptic curve (C, x_0) and recall that the Weil pairing on (C, x_0) is defined in terms of such functions (see [Sil, III.8]). By Riemann-Roch the dimension of $\mathcal{L}(\pi^*[x_0])$ is $p^2 = \#X$. Clearly the G_x lie in the Riemann-Roch space, so it will suffice to show that they are linearly independent. This follows from the definition of the Weil pairing; the G_x are eigenfunctions for distinct characters with respect to the action of $X = C[p]$ by translation. (see the first paragraph of the proof of Prop. 3.3 in [CFOSS-II]).

Thus we may define an embedding of C into \mathbb{P}^{p^2-1} via

$$g : C \ni P \mapsto (G_x(P))_{x \in X} \in \mathbb{P}^{p^2-1}.$$

It is evident that the pull-back of any flex point $x \in X$ by π is the hyperplane section of $g(C) \subset \mathbb{P}^{p^2-1}$ cut out by $z_x = 0$. Let $Q \in C \setminus C[p^2]$ be any point, with projective coordinates $g(Q)$. By the definition of the descent map, the image of (C, π) under the descent map is represented by the $\Delta \in \bar{H}^\times$ such that

$$(7.6) \quad \tilde{\ell}(\pi(Q)) = \Delta \partial(g(Q)).$$

By definition we have that $\Delta \in \tilde{\mathcal{H}}_{\bar{K}}$.

Equation (7.6) was also used to construct D_Δ . So it is clear that $\pi_\Delta \circ g = \pi$ on $C \setminus C[p^2]$ and that the image of this open subscheme under g is contained in D_Δ . Since D_Δ is projective (hence complete), this is then true on all of C . We conclude that $g(C) \subset D_\Delta$ and that $\pi_\Delta \circ g = \pi$. On the other hand, $g(C)$ is a genus one normal curve of degree p^2 . Its homogeneous ideal can be generated by a \bar{K} -vector space of quadrics of dimension $p^2(p^2-3)/2$. We have already found a set of $p^2(p^2-3)/2$ linearly independent quadrics vanishing on D_Δ in lemma 7.3, so we must have $g(C) = D_\Delta$. Thus (D_Δ, π_Δ) is a twist of (C, π) . This completes the proof. \square

Alternate proof of Theorem 7.1. The second proof involves a more direct examination of the defining equation(s) of D_Δ . Again we work over \bar{K} with the basis given by characteristic functions. We will also suppress the base field from the notation, writing $\tilde{\mathcal{H}}$ and \mathcal{H} for $\tilde{\mathcal{H}}_{\bar{K}}$ and $\mathcal{H}_{\bar{K}}$.

The linear form $\tilde{\ell}$ splits completely over \bar{H} as $\tilde{\ell} = (\tilde{\ell}_y)_{y \in Y}$, the $\tilde{\ell}_y$ being linear forms defining the hyperplanes $y \in Y$. The equation

$$\tilde{\ell} = a\Delta\partial(z)$$

used to define D_Δ becomes the system of equations (indexed by $y \in Y$)

$$\tilde{\ell}_y(u_1, \dots, u_p) = a\Delta_y \prod_{x \in y} z_x,$$

where as usual the product is to be taken with appropriate multiplicities. We will also use the index x to denote the factor of \bar{H} corresponding to the hyperplane in $Y_1 \subset Y$ which cuts out the divisor $p[x]$ on C . Recall also that in this basis, the action of \bar{F}^\times on \mathbb{P}^{p^2-1} is simply coordinate-wise multiplication.

Construction of D_δ . Let $\Delta \in \bar{H}^\times$. Under the splitting $\bar{H} \simeq \bar{F} \times \bar{H}_2$, we will write $\Delta = (\delta, \varepsilon)$. In terms of the basis over \bar{K} , $\delta = (\delta_x)_{x \in X}$ and $\varepsilon = (\varepsilon_y)_{y \in Y_2}$. The p^2 hyperplane sections of C supported on a single point (i.e. the $p[x]$ for $x \in X$), correspond to equations

$$\tilde{\ell}_x = a\delta_x z_x^p, \quad x \in X.$$

In completely analogous fashion to the construction of D_Δ at the beginning of this section, one can use these p^2 equations, together with the homogeneous polynomials defining $C \subset \mathbb{P}^{p^2-1}$, to define a C -scheme, $D_\delta \subset \mathbb{P}^{p^2-1}$.

For $p = 3$, D_δ is given by 6 cubic forms and one form of degree 9. For larger p it is given by $p^2 - p$ forms of degree p and $p(p - 3)/2$ forms of degree $2p$. In both cases the map $\pi_\delta : D_\delta \rightarrow C$ is given by p forms of degree p . The construction works for any $\delta = (\delta_x) \in \bar{F}^\times$ and (D_δ, π_δ) can be defined over the minimal field of definition of δ and only depends on the class of δ modulo K^\times .

LEMMA 7.7. *D_δ is a curve³ in \mathbb{P}^{p^2-1} . The map $\pi_\delta : D_\delta \rightarrow C$ is of degree p^{p^2-1} . It ramifies only above the flex points of C where the ramification index is p .*

PROOF: For all of these statements it suffices to count preimages of points $P \in C$. The points of D_δ above $P \in C$ correspond to solutions to the system of equations $\tilde{\ell}_x(P) = \delta_x z_x^p$. Note that, since $\tilde{\ell}_x$ defines a hyperplane meeting C only at the point x , $\tilde{\ell}_x(P) = 0$ if and only if $P = x$.

For $P \notin X$ we get p choices for each z_x (the p -th roots of $\tilde{\ell}_x(P)/\delta_x$), which in homogeneous coordinates correspond to p^{p^2-1} points on D_δ . For $P \in X$ exactly one of the $\tilde{\ell}_x(P)$ is 0, so we have a factor of p fewer solutions. \square

The splitting of D_δ . If $\Delta = (\delta, \varepsilon) \in \bar{F}^\times \times \bar{H}_2^\times$, then D_Δ is a (possibly empty) subscheme of D_δ . The following lemma shows that it is those $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$ that provide something interesting.

LEMMA 7.8. *Let $(\delta, \varepsilon) \in \bar{H}^\times$. If $(\delta, \varepsilon) \notin \tilde{\mathcal{H}}$, then $D_{(\delta, \varepsilon)}$ consists of finitely many points lying above the flex points of C . If $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$, then any $Q \in D_\delta$ is contained in $D_{(\delta, \varepsilon')}$ for some $(\delta, \varepsilon') \in \tilde{\mathcal{H}}$.*

PROOF: Suppose $Q \in D_\delta$ lies over $P \notin X$. If $Q \in D_{(\delta, \varepsilon)}$, then (for an appropriate choice of homogeneous coordinates for P and Q) we have $\tilde{\ell}(P) = (\delta, \varepsilon)\partial(Q)$. By definition, this implies $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$, which proves the first statement.

If (δ, ε) is in $\tilde{\mathcal{H}}$, we can write $(\delta, \varepsilon) = (d, e) \cdot \partial(\alpha)$ where $(d, e) = \tilde{\ell}(P)$ and $\alpha \in \bar{F}^\times$. Then for all $x \in X$, $\tilde{\ell}_x(P) = \delta_x Q_x^p = d_x (\alpha_x Q_x)^p$. The divisors on C cut out by $\partial(\tilde{\ell}_x)_y$ and $(\tilde{\ell}_y)^p$ are equal, so there are constants $\beta_y \in \bar{K}^\times$ (see lemma 5.5) such that for all $y \in Y_2$, $\partial_2(\tilde{\ell}_x(P))_y = \beta_y \tilde{\ell}_y(P)^p$ and $\partial_2(d_x)_y = \beta_y e_y^p$.

Applying ∂_2 to the equations $\tilde{\ell}_x(P) = d_x (\alpha_x Q_x)^p$ gives:

$$\beta_y \tilde{\ell}_y(P)^p = \beta_y e_y^p \partial_2(\alpha_x Q_x)_y^p, \quad \text{for } y \in Y_2.$$

Canceling β_y and taking a p -th root of this relation, we see that there is some $e' = (e'_y) \in \bar{H}_2^\times$ such that

$$\tilde{\ell}_y(P) = e'_y \partial(\alpha_x Q_x)_y.$$

³We are not claiming D_δ is irreducible nor that it is smooth.

From this it is clear that $Q \in D_{(d,e') \cdot \partial(\alpha)}$. Since Q does not lie above a flex, the first part of the lemma implies that $(d, e') \cdot \partial(\alpha) \in \tilde{\mathcal{H}}$.

This proves the second statement under the assumption that Q does not lie above any flex point. The statement for the finitely many remaining Q is given in the proof of 7.14 below. One can check that the logical dependence is not circular. \square

Over \bar{K} , the fibers of the projection $\text{pr}_1 : \mathcal{H} \rightarrow \bar{F}^\times / \bar{K}^\times \bar{F}^{\times p}$ are trivial (since both groups are in fact trivial). At the level of representatives we have the following.

LEMMA 7.9. *Let $\delta \in \bar{F}^\times$. The preimage of δ in $\tilde{\mathcal{H}}$ is either empty, or is a coset of $\partial(\mu_p(\bar{F})) \subset \bar{H}^\times$. In particular, the number of preimages is either 0 or p^{p^2-3} .*

PROOF: We have an exact sequence

$$\begin{array}{ccccccc} 1 & \longrightarrow & \partial(\mu_p(\bar{F})) & \longrightarrow & \partial\bar{F}^\times & \xrightarrow{\text{pr}_1} & \bar{F}^\times & \longrightarrow & 1 \\ & & & & \downarrow & & \parallel & & \\ & & & & \bar{H}^\times & \xrightarrow{\text{pr}_1} & \bar{H}_1^\times & & \end{array}$$

where the map on the right is projection onto the first factor. But $\tilde{\mathcal{H}}$ is a coset of $\partial\bar{F}^\times = \tilde{\mathcal{H}}^0$, so the first statement is clear.

For the second statement, recall that ∂ was chosen so that its kernel is the \mathbb{F}_p -vector space of affine maps $\text{Aff}(X, \mu_p)$. Hence

$$\dim_{\mathbb{F}_p} \partial(\mu_p(\bar{F})) = \dim_{\mathbb{F}_p} \mu_p(\bar{F}) - \dim_{\mathbb{F}_p} \text{Aff}(X, \mu_p) = p^2 - 3,$$

as required. \square

The numerology suggests that D_δ , which is of degree p^{p^2-1} over C , should split into $p^{p^2-3} = \#\partial(\mu_p(\bar{F}))$ components each of degree p^2 over C and corresponding to some (δ, ε) in the fiber over δ . For this to work out, we first need to check that the $D_{(\delta, \varepsilon)}$ are in fact distinct. For this we make use of the action of $\mu_p(\bar{F})/\mu_p$. Recall that $\bar{F}^\times / \bar{K}^\times$ acts on \mathbb{P}^{p^2-1} by automorphisms. In terms of the basis given by the characteristic functions, this is just coordinate-wise multiplication.

Given $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$, the fiber over δ is $\{\partial(\eta) \cdot (\delta, \varepsilon) : \eta \in \mu_p(\bar{F})\} \subset \tilde{\mathcal{H}}$. By lemma 7.4, the corresponding C -schemes, $D_{\partial(\eta) \cdot (\delta, \varepsilon)}$, are permuted transitively by the action of $\mu_p(\bar{F})/\mu_p \subset \bar{F}^\times / \bar{K}^\times$ on \mathbb{P}^{p^2-1} .

LEMMA 7.10. *Under the action of $\mu_p(\bar{F})/\mu_p$, $\eta \cdot D_{(\delta, \varepsilon)} = D_{(\delta, \varepsilon)}$ if and only if $\partial(\eta) = 1$. If $\eta \cdot D_{(\delta, \varepsilon)} \neq D_{(\delta, \varepsilon)}$, then $\eta \cdot D_{(\delta, \varepsilon)} \cap D_{(\delta, \varepsilon)}$ is contained in the union of the hyperplanes $\{z_x = 0\}$ for $x \in X$.*

PROOF: Suppose $Q \in D_{(\delta, \varepsilon)}$ with homogeneous coordinates $(Q_x)_{x \in X}$. In order that $Q \in D_{\partial(\eta) \cdot (\delta, \varepsilon)}$ it is necessary and sufficient that

$$\varepsilon_y \prod_{x \in y} Q_x = \partial(\eta)_y \varepsilon_y \prod_{x \in y} Q_x = \varepsilon_y \prod_{x \in y} \eta_x Q_x$$

for each $y \in Y$ (up to a constant in \bar{K}^\times , uniform in y). This is equivalent to requiring that $\partial(\eta)$ take the same value at all $y \in Y$ such that $\prod_{x \in y} Q_x \neq 0$.

If Q does not lie above a flex point, $\prod_{x \in y} Q_x$ is nonzero for all y , and so the above holds if and only if $\partial(\eta)$ is constant. Since $\partial_1(\eta) = \eta^p = 1$, this implies $\partial(\eta) = 1$.

This proves the first statement. The second statement follows from the fact that these hyperplanes contain the fibers above the flex points. \square

DEFINITION 7.11. *We call a class in $\mu_p(\bar{F})/\mu_p$ a pseudo-reflection about the hyperplane $\{z_x = 0\}$ if it is represented by a map $\eta \in \mu_p(\bar{F}) = \text{Map}(X, \mu_p)$ with value equal to 1 at all flex points not equal to x , and equal to a nontrivial p -th root of unity at x .*

REMARK: The terminology comes from the fact that the action of a pseudo-reflection about $\{z_x = 0\}$ on \mathbb{P}^{p^2-1} leaves the hyperplane fixed and has order p . A pseudo-reflection of order 2 is a reflection.

LEMMA 7.12. *$\mu_p(\bar{F})/\mu_p$ acts on D_δ by automorphisms (as a C -scheme). A point $Q \in D_\delta$ is fixed by $\eta \in \mu_p(\bar{F})/\mu_p$ if and only if Q lies above the flex point x and η is a pseudo-reflection about the hyperplane $\{z_x = 0\}$.*

PROOF: The fact that $\mu_p(\bar{F})/\mu_p$ acts by automorphisms on D_δ is clear from the defining equations, $\tilde{\ell}_x = \delta_x z_x^p$.

For the second statement, let $Q = (Q_x) \in D$. If η represents some class $\tilde{\eta} \in \mu_p(\bar{F})/\mu_p$, then $\tilde{\eta}Q = Q$ if and only if there is some $c \in \bar{K}^\times$ such that $cQ_x = \eta_x Q_x$ for all $x \in X$. This happens if and only if the value of η is the same at all x where $Q_x \neq 0$. But $Q_x = 0$ if and only if Q is in the fiber above $x \in X$. So the statement is clear. \square

COROLLARY 7.13. *If $P \in C$ is any point, then $\mu_p(\bar{F})/\mu_p$ acts transitively on the fiber $\pi_\delta^{-1}(P)$.*

PROOF: The stabilizer S_P of any fiber is either trivial (if $P \notin X$) or of order p (if $P \in X$). Comparing with 7.7 we see that in either case

$$\#\pi_\delta^{-1}(P) \cdot \#S_P = \#\mu_p(\bar{F})/\mu_p.$$

So the action is transitive. \square

PROPOSITION 7.14. *Let $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$. Then D_δ splits as a union of distinct C -schemes*

$$D_\delta = \bigcup D_{(\delta, \varepsilon)},$$

the union running over the p^{p^2-3} elements $(\delta, \varepsilon) \in \tilde{\mathcal{H}}$ in the fiber above δ . Moreover, each $D_{(\delta, \varepsilon)}$ is an unramified covering of C of degree p^2 .

We prove the proposition below. First we give two important corollaries and some discussion of how this relates to the descent map.

COROLLARY 7.15. *Suppose $\Delta \in \tilde{\mathcal{H}}_K$. If D_Δ is connected, then (D_Δ, π_Δ) is a p -covering of C representing a class in $\text{Cov}_0^{(p)}(C/k)$ and its image under the descent map is the class of Δ .*

PROOF: We know that $\pi_\Delta : D_\Delta \rightarrow C$ is an unramified covering of degree p^2 defined over K . Lemma 7.10 shows that

$$E[p] \simeq \ker(\mu_p(\bar{F})/\mu_p \xrightarrow{\partial_2} \mu_p(\bar{H}_2))$$

acts on D_Δ by automorphisms which are simply transitive on each fiber. It follows that (D_Δ, π_Δ) represents a class in $\text{Cov}^{(p)}(C/K)$. Since we have exhibited a model of degree p^2 in \mathbb{P}^{p^2-1} , this class is in $\text{Cov}_0^{(p)}(C/K)$. By construction, the image of this class under the descent map is the class of Δ in \mathcal{H}_K (cf. 7.2). \square

Assume δ is as in the proposition and in addition that $\delta \in F^\times$. Then D_δ is defined over K . The splitting in the proposition occurs over \bar{K} . The set of components defined over K is in one to one correspondence with the preimage of δ in $\tilde{\mathcal{H}}_K$. This preimage is either empty (in which case δ does not represent the image of any p -covering under the fake descent map) or is a coset of $(\partial_2(\mu_p(\bar{F})))^{G_K}$ (cf. lemma 7.9 and the diagrams of section 6). The G_K -invariant subgroup of $\mu_p(\bar{F})/\mu_p$ acts on these K -defined components, its orbits being the K -isomorphism classes. This gives a geometric interpretation of the group \mathcal{K}_K appearing in the diagrams of section 6. Namely it parametrizes the K -isomorphism classes of the K -defined components of D_δ under the splitting given in the proposition. These components, in turn, represent the isomorphism classes of p -coverings of C defined over K whose image under the fake descent map is equal to the class of δ .

We have seen in lemma 7.12 that $\text{Aut}_C(D_\delta) \simeq \mu_p(\bar{F})/\mu_p$. By the twisting principle, the twists of D_δ are parameterized by the group $H^1(K, \mu_p(\bar{F})/\mu_p)$ appearing in the diagrams of section 6. Using the explicit description of the action $\mu_p(\bar{F})/\mu_p$ on \mathbb{P}^{p^2-1} we can write down matrices in PGL_{p^2} giving the action of $E[p]$ on $D_{(\delta, \varepsilon)}$. This will likely be useful when trying to perform minimization and reduction to obtain ‘nicer’ models.

As a second corollary, we have a proof of lemma 5.2 above.

COROLLARY 7.16. *Suppose $\Delta \in H^\times$ and $K \subset L$ is any extension of fields. If $\Delta \otimes 1_L \in \tilde{\mathcal{H}}_L$, then $\Delta \in \tilde{\mathcal{H}}_K$.*

PROOF: Lemma 7.8 and the proposition show that D_Δ has positive dimension if and only if $\Delta \in \tilde{\mathcal{H}}_K$. But this is a geometric property. In particular, $D_{\Delta \otimes 1_L} = D_\Delta \times_{\text{Spec}(K)} \text{Spec}(L)$ can only be of positive dimension if D_Δ is as well. \square

Now we prove the proposition.

PROOF OF PROPOSITION 7.14: That D_δ splits as in the statement follows from lemma 7.8, but the proof of this lemma is not complete. It remains to show that points in the fiber above a flex point lie on some $D_{(\delta, \varepsilon)}$ in the union. For this note that each $D_{(\delta, \varepsilon)}$ appearing is a curve with a non-constant hence surjective map to C . Hence, given $x \in X$, there exists some $Q \in D_{(\delta, \varepsilon)}$ above x . By 7.13, the fiber above x in D_δ is the orbit of x under the action of $\mu_p(\bar{F})/\mu_p$. But each point in this orbit lies on some $D_{(\delta, \varepsilon)}$ appearing in the union.

We have seen in lemma 7.9 that the union runs over a set of size p^{p^2-3} . If $P \notin X$, then the fibers above P on distinct $D_{(\delta, \varepsilon)}$ are disjoint (7.10). The action of $\mu_p(\bar{F})/\mu_p$ then shows that fiber above P on any $D_{(\delta, \varepsilon)}$ must have size $p^{p^2-1}/p^{p^2-3} = p^2$. So each $D_{(\delta, \varepsilon)}$ is of degree p^2 over C .

The only possible branch points are the flex points, since these are the only branch points of D_δ . On D_δ all the ramification points have index p . So it suffices to show that at least p distinct $D_{(\delta, \varepsilon)}$ come together at each $Q \in D_\delta$ lying above a flex. It is easy to see that if η is a pseudo-reflection about $\{z_x = 0\}$, then $\partial(\eta) \neq 1$. Now using lemmas

7.10 and 7.12, this shows that the orbit of any $D_{(\delta,\varepsilon)}$ under η consists of p distinct curves, which all share a common fiber over x . This completes the proof. \square

In order to complete the proof of Theorem 7.1 (without using lemma 7.6) one would need to show that the curve D_Δ associated to some $\Delta \in \tilde{\mathcal{H}}_K$ is (geometrically) connected. Let us suppose for the moment that D_Δ splits (over \bar{K}) into n irreducible components D_i . Since D_Δ is non-singular, the D_i are necessarily disjoint inside \mathbb{P}^{p^2-1} . The action of $E[p] \subset \mu_p(\bar{F})/\mu_p$ must permute the components transitively, so there are three possibilities: $n \in \{1, p, p^2\}$. If $n = 1$, then D_Δ is connected. If $n = p^2$, then as D_Δ is a curve of degree p^2 , we see that each irreducible component is a genus one curve of degree one in \mathbb{P}^{p^2-1} . This is absurd since it would imply the existence of a very ample divisor of degree one on a curve of genus one.

The case remaining to be ruled out is that D_Δ splits as a union of p genus one curves of degree p . If this is the case, then each D_i is a Galois covering of C with Galois group isomorphic to a cyclic subgroup of $E[p]$. It is not difficult to see that in fact all p of these coverings have the same automorphism group and, consequently, that this subgroup is defined over K (as a whole, not necessarily point-wise). In general, $E[p]$ need not contain any K -rational cyclic subgroup, so ‘generically’ this case does not occur. Nevertheless, it seems difficult to turn this into a proof (without resorting to something along the lines of 7.6).

CHAPTER III

Computing the p -Selmer Set

We shift our focus now to the arithmetic situation. We specialize to the case that $K = k$ is a number field. We assume that C is an everywhere locally solvable genus one normal curve of degree p defined over k , with a fixed model as described in the beginning of the previous chapter. Recall that local solvability implies that $\text{Pic}(C) = \text{Pic}_k(C)$ and that $\text{Cov}^{(p)}(C/k) \neq \emptyset$. Thus all of the material of the previous section applies to C (see the discussion at the beginning of Chapter II).

An element in an étale k -algebra $A \simeq \prod_i K_i$ will be called *integral* if its image in each K_i is integral. We assume that the linear form $\tilde{\ell}$ defining the descent map and all polynomials appearing in the model for C have integral coefficients. We further assume that the constants $c \in k^\times$ and $\beta \in H_2^\times$ given by II.5.5 are integral. All of this can be achieved by scaling. We denote the completion of k at a prime v by k_v . We attach a subscript v to any object defined over k to denote the corresponding object over k_v obtained by extension of scalars. For example $H_v = H \otimes k_v$, $\tilde{\mathcal{H}}_v = \tilde{\mathcal{H}}_{k_v}$, $\text{Pic}_v(C) = \text{Pic}_{k_v}(C)$, and so on.

1. The algebraic Selmer set

The descent map allows us to identify $\text{Sel}^{(p)}(C/k)$ with its image in $H^\times/k^\times \partial F^\times$. We now determine the image. This gives an algebraic presentation of the p -Selmer set, which can be computed fairly directly.

Consider the following diagram:

$$\begin{array}{ccc} \text{Pic}_k(C) & \xrightarrow{\Phi} & H^\times/k^\times \partial F^\times \\ \downarrow & & \downarrow \Pi \text{res}_v \\ \prod_v \text{Pic}_v(C) & \xrightarrow{\Pi \Phi_v} & \prod_v H_v^\times/k_v^\times \partial F_v^\times \end{array}$$

If $(D, \pi) \in \text{Sel}^{(p)}(C/k)$ is an everywhere locally solvable p -covering of C , then its image, $\tilde{\Phi}((D, \pi)) \in H^\times/k^\times \partial F^\times$, has the property that it maps under $\prod_v \text{res}_v$ into the subset $\prod_v \Phi_v(\text{Pic}_v^1(C)) \subset \prod_v \mathcal{H}_v$. This suggests the following definition.

DEFINITION 1.1. *The algebraic p -Selmer set of C associated to Φ is the set*

$$\text{Sel}_{alg}^{(p)}(C/k) = \{ \Delta \in H^\times/k^\times \partial F^\times : \text{res}_v(\Delta) \in \Phi_v(\text{Pic}_v^1(C)) \text{ for all } v \}.$$

THEOREM 1.2. *The descent map gives a one to one correspondence between the p -Selmer set of C and the algebraic p -Selmer set of C .*

PROOF: The defining property of the descent map shows that the image of $\text{Sel}^{(p)}(C/k)$ is equal to $\text{Sel}_{alg}^{(p)}(C/k) \cap \mathcal{H}_k$. We know that the descent map is injective by II.4.2, so

it suffices to show that $\text{Sel}_{alg}^{(p)}(C/k) \subset \mathcal{H}_k$. This follows from II.5.2; any element of $H^\times/K^\times \partial F^\times$ which restricts into \mathcal{H}_v (for some v) is an element of \mathcal{H}_k . \square

One can formulate the same definition for divisor classes of degree 0.

DEFINITION 1.3. *The algebraic p -Selmer group of $E = \text{Jac}(C)$ is*

$$\text{Sel}_{alg}^{(p)}(E/k) = \{ \Delta \in H^\times/k^\times \partial F^\times : \text{res}_v(\Delta) \in \Phi_v(\text{Pic}_v^0(C)) \text{ for all } v \}.$$

Note that by II.5.2, $\text{Sel}_{alg}^{(p)}(E/k) \subset \mathcal{H}_k^0$. Since \mathcal{H}_k is a principal homogeneous space for \mathcal{H}_k^0 , the same is true of the corresponding Selmer objects.

LEMMA 1.4. *If the algebraic p -Selmer set of C is nonempty, then it is a coset of the algebraic p -Selmer group of E inside $H^\times/k^\times \partial F^\times$.*

PROOF: By assumption C is everywhere locally solvable. So everywhere locally, the group of k_v -rational divisor classes of degree 1 on C is a coset of the group of k_v -rational divisor classes of degree 0. Since Φ_v is a homomorphism, the same is true of their images in $H_v^\times/k_v^\times \partial F_v^\times$. If the algebraic p -Selmer set of C is nonempty, then these cosets can be simultaneously defined by some global element of \mathcal{H}_k . \square

In section II.1 we defined the fake Selmer set. If it is empty, then so is the p -Selmer set. Using the previous lemma, one can do slightly better. The group \mathcal{K}_k was defined in section II.6.

COROLLARY 1.5. *If $\#\text{Sel}_{fake}^{(p)}(C/k) < \frac{\#\text{Sel}^{(p)}(E/k)}{\#\mathcal{K}_k}$, then $\text{Sel}^{(p)}(C/k) = \emptyset$.*

PROOF: The projection $H^\times \xrightarrow{\text{pr}_1} F^\times \times H_2^\times$ induces maps

$$\begin{array}{ccc} \text{Sel}_{alg}^{(p)}(C/k) & \xrightarrow{\text{pr}_1} & \text{Sel}_{fake}^{(p)}(C/k) \\ \downarrow & & \downarrow \\ \mathcal{H}_k^0 & \xrightarrow{\text{pr}_1} & F^\times/K^\times F^{\times p} \end{array}$$

where the kernel of the lower map is \mathcal{K}_k . This gives a bound on the size of the fibers of the upper map. On the other hand, the lemma shows that $\#\text{Sel}_{alg}^{(p)}(C/k)$ is equal to 0 or $\#\text{Sel}^{(p)}(E/k)$. \square

Although it is not reflected in the notation, $\text{Sel}_{alg}^{(p)}(C/k)$ depends on our choice of linear form $\tilde{\ell}$ used to define the descent map and the algebraic p -Selmer group of E depends on C . The next proposition shows, however, that the image of $\text{Sel}_{alg}^{(p)}(E/k)$ in $H^1(k, E[p])$ does not.

PROPOSITION 1.6. *The inclusion $\mathcal{H}_k^0 \simeq C^\perp \hookrightarrow H^1(k, E[p])$ identifies the algebraic p -Selmer group of E with the p -Selmer group of E .*

PROOF: We identify \mathcal{H}_k^0 with its image in $H^1(k, E[p])$ and E with $\text{Pic}_k^0(C)$. To show that the algebraic Selmer group is contained in the Selmer group we use lemma II.4.5. This says that the images of $\Phi_v|_{\text{Pic}_v^0(C)}$ and the connecting homomorphism δ_v from the Kummer sequence of E/k_v are the same. So clearly the algebraic Selmer group is contained in the Selmer group.

For the reverse inclusion it suffices to show that $\text{Sel}^{(p)}(E/k) \subset \mathcal{H}_k^0 \simeq C^\perp$. So we need to show that elements of the Selmer group are orthogonal to C with respect to the Weil pairing induced cup product of level p . Using that the cup product is the bilinear form associated to the obstruction map we have

$$C \cup_p C' = \text{Ob}_p(C + C') - \text{Ob}_p(C) - \text{Ob}_p(C'),$$

for any $C' \in H^1(k, E[p])$. If C' is everywhere locally solvable then so is $C + C'$ (because the Selmer group is a group). Having points everywhere locally implies trivial obstruction, so all the terms on the right-hand side vanish as required. \square

Computable description. In order to compute the algebraic Selmer set explicitly, we need a method for determining these local images. For a given v , this is relatively straight-forward (see section 2). But there are infinitely many primes to deal with. As is the case for p -descent on elliptic curves, algebraic number theory gives us a means for handling all but finitely many ‘bad primes’ simultaneously.

For a completion k_v of k at a non-archimedean prime, we use k_v^{unr} to denote the maximal unramified extension of k_v . If ξ is an element of some object defined over k , we say that ξ is unramified at v if ξ becomes trivial upon extension of scalars to k_v^{unr} . For Galois cohomology groups $H^1(k, -)$, this coincides with the usual definition that ξ be in the kernel of the restriction map to $H^1(k_v^{\text{unr}}, -)$. For example, a class in $F^\times/k^\times F^{\times p}$ represented by δ is unramified at v if $\delta \in k_v^{\text{unr}\times}(F \otimes k_v^{\text{unr}})^{\times p}$ or, equivalently if its image under the map $F^\times/k^\times F^{\times p} \hookrightarrow H^1(k, \mu_p(\bar{F})/\mu_p) \rightarrow H^1(k_v^{\text{unr}}, \mu_p(\bar{F}_v)/\mu_p)$ is zero.

The first step is to identify a suitable finite set of ‘bad primes’. To that end, let F' denote the field extension of k obtained by adjoining the coordinates of all flex points of C . We refer to F' as the splitting field of X . In case F is a field (i.e. the action of G_k on X is transitive), F' is the normal closure of F . In general¹, we can write F as a quotient of the polynomial ring $k[\tau]$ by some polynomial $f(\tau)$ for which F' is the splitting field.

Write the linear form used to define the descent map as $\tilde{\ell} = (\tilde{\ell}_1, \tilde{\ell}_2)$ under the splitting $H \simeq F \times H_2$. Here $\tilde{\ell}_1$ defines a hyperplane section meeting C at a generic flex point with multiplicity p . Over F' , all flex points are defined, and so $\tilde{\ell}_1$ splits as a p^2 -tuple, $(\tilde{\ell}_x)_{x \in X}$ of linear forms with coefficients in F' each defining the hyperplane meeting C only at the flex $x \in X$.

At any non-archimedean prime w of F' , we can reduce the $\tilde{\ell}_x \bmod w$. Since this linear form may vanish $\bmod w$, it may fail to define a hyperplane section of the reduction of $C \bmod w$. In some sense this is a situation we would like to avoid. We will refer to a prime v of k as a prime of bad reduction (resp. good reduction) for $\tilde{\ell}$ if there is some (resp. no) prime $w|v$ of F' for which this occurs.

REMARK: Even when working over \mathbb{Q} it may not be possible to choose $\tilde{\ell}$ in such a way that it has good reduction everywhere, since F can have nontrivial class group. In any event, such bad primes can be detected quite easily.

¹i.e. over any infinite field. This is not true of every étale algebra over a finite field.

Recall the constant $c \in K^\times$ from lemma II.5.5 defined, up to p -th powers, by the property that $N_{F/k}(\delta) \equiv c \pmod{k^{\times p}}$ for any $(\delta, \varepsilon) \in \mathcal{H}_k$. We may scale to ensure that c is integral.

LEMMA 1.7. *Let v be a non-archimedean prime of k which is of good reduction for both C and $\tilde{\ell}$ and which is prime to both p and c . Then $\Phi_v(\text{Pic}_v^1(C)) \subset H_v^\times/k_v^\times \partial F_v^\times$ is contained in the unramified subgroup.*

PROOF: Let F' be the splitting field of X . By the criterion of Neron-Ogg-Shafarevich, the primes which ramify in the extension F'/k are either primes of bad reduction for C or lie above p . In particular, if v is as in the lemma, then it does not ramify in F' .

Now we claim that if v does not ramify in F' , then for all $(\delta, \varepsilon) \in \mathcal{H}_v$ we have

$$(\delta, \varepsilon) \in \mathcal{H}_v \text{ is unramified} \Leftrightarrow \delta \in F_v^\times/K_v^\times F_v^{\times p} \text{ is unramified.}$$

This follows from the fact that for these ‘good primes’ the map

$$\mathcal{H}_{k_v^{\text{unr}}} \rightarrow F^{\text{unr}\times}/K^{\text{unr}\times} F^{\text{unr}\times p}$$

induced by projection onto the first factor of $H \simeq F \times H_2$ is injective. To see the injectivity recall that the fibers of this map (see the diagrams in II.6) are parameterized by

$$\mathcal{K}_v := \frac{H^0(k_v^{\text{unr}}, (\partial_2(\mu_p \bar{F})))}{\partial_2 \left(H^0 \left(k_v^{\text{unr}}, \frac{\mu_p(\bar{F})}{\mu_p} \right) \right)}.$$

As v does not ramify in F' , all flex points are defined over k_v^{unr} , so the action on the modules appearing here is trivial. Since $\mu_p \subset \ker \partial_2$, we have $\partial_2(\mu_p(\bar{F})) = \partial_2(\mu_p(\bar{F})/\mu_p)$. So the quotient is trivial.

To prove the lemma, it now suffices to show that the image of the composition

$$\Phi_{\text{fake},v} : \text{Pic}_v^1(C) \xrightarrow{\Phi_v} \mathcal{H}_v \xrightarrow{\text{pr}_1} F_v^\times/k_v^\times F_v^{\times p}$$

is unramified. For this it will be enough to show that this is true of any point $P \in C(k_v)$ which is neither a zero nor a pole of ℓ_1 . For this we can choose primitive integral coordinates for P (i.e. homogeneous coordinates with valuations that are non-negative but not all positive) and consider $\tilde{\ell}_1(P) \in (F \otimes k_v)^\times$. The flex algebra $F \otimes k_v$ splits as a product of extensions of k_v . Since $v \nmid p$, in order that the image of P be unramified it is sufficient that the valuation of $\tilde{\ell}_1(P)$ in each of these factors is a multiple of p .

Fix some factor $K_{\mathfrak{v}}$. For any prime \mathfrak{w} of F' extending \mathfrak{v} , we get an unramified tower of fields $k_{\mathfrak{v}} \subset K_{\mathfrak{v}} \subset F'_{\mathfrak{w}}$. Let $\nu_{\mathfrak{w}}$ be the normalized valuation on $F'_{\mathfrak{w}}$. Over F' , $\tilde{\ell}_1$ splits as $(\tilde{\ell}_x)_{x \in X}$ and, since the extensions are all unramified, it suffices to show that $\nu_{\mathfrak{w}}(\tilde{\ell}_x(P)) \equiv 0 \pmod{p}$ for each $x \in X$.

For this we make use of the norm condition. Since $v \nmid c$, its valuation satisfies the congruence $\nu_{\mathfrak{w}}(c) \equiv 0 \pmod{p}$. Hence,

$$\sum_{x \in X} \nu_{\mathfrak{w}}(\tilde{\ell}_x(P)) = \nu_{\mathfrak{w}} \left(\prod_{x \in X} \tilde{\ell}_x(P) \right) \equiv \nu_{\mathfrak{w}}(c) \equiv 0 \pmod{p}.$$

Each summand appearing on the left is nonnegative. To complete the proof it suffices to show that at most one can be positive.

Since v is of good reduction for $\tilde{\ell}$, the reduction of each $\tilde{\ell}_x$ defines a hyperplane meeting \tilde{C} only at the image \tilde{x} of x on \tilde{C} . So $\nu_{\mathfrak{w}}(\tilde{\ell}_x(P)) > 0$ if and only if P and x have the same image under the reduction map. On the other hand, the images of the flex points modulo \mathfrak{w} are all distinct since v is of good reduction for C and is prime to p .

(the images of these flex points are the flex points of the reduced curve). So $\nu_{\mathfrak{w}}(\tilde{\ell}_x(P))$ can be positive for at most one $x \in X$. This completes the proof. \square

PROPOSITION 1.8. *Let S be the set of primes of k containing all non-archimedean primes dividing p or c , all primes of bad reduction for C or $\tilde{\ell}$ and all archimedean primes if $p = 2$. Let \mathcal{H}_S denote the subgroup of \mathcal{H}_k consisting of elements that are unramified outside S . Then*

$$\mathrm{Sel}_{\mathrm{alg}}^{(p)}(C/k) = \{ \Delta \in \mathcal{H}_S : \mathrm{res}_v(\Delta) \in \Phi_v(\mathrm{Pic}_v^1(C)), \text{ for all } v \in S \}.$$

PROOF: Let Z denote the set in the statement. The previous lemma shows that Z contains $\mathrm{Sel}_{\mathrm{alg}}^{(p)}(C/k)$. To show that the reverse inclusion holds, we may assume that Z is nonempty. Let $\Delta \in Z$ and $v \notin S$. To show that $\Delta \in \mathrm{Sel}_{\mathrm{alg}}^{(p)}(C/k)$ we must show that $\mathrm{res}_v(\Delta) \in \Phi_v(\mathrm{Pic}_v^1(C))$. Choose any $P \in \mathrm{Pic}_v^1(C)$. Then both $\Phi_v(P)$ and $\mathrm{res}_v(\Delta)$ are unramified, so $\Phi_v(P) \cdot \mathrm{res}_v(\Delta)^{-1}$ is in the unramified subgroup of \mathcal{H}_v^0 .

Since v is a prime of good reduction for C , it is also a prime of good reduction for its Jacobian. By I.4.7, the image of the connecting homomorphism $\delta_v : E(k_v) \rightarrow H^1(k_v, E[p])$ is equal to the unramified subgroup. On the other hand, II.4.5 says that $\Phi_v : \mathrm{Pic}_v^0(C) \rightarrow \mathcal{H}_v^0 \subset H^1(k_v, E[p])$ coincides with connecting homomorphism. It follows that $\Phi_v(\mathrm{Pic}_v^0(C))$ is equal to the unramified subgroup of \mathcal{H}_v^0 . Hence there exists some $Q \in \mathrm{Pic}_v^0(C)$ such that $\Phi_v(Q) = \Phi_v(P) \mathrm{res}_v(\Delta)^{-1}$. Since Φ_v is a homomorphism we have $\mathrm{res}_v(\Delta) = \Phi_v(P - Q)$, which completes the proof since $P - Q \in \mathrm{Pic}_v^1(C)$. \square

REMARK: Let S be as in the proposition and let S_1 be the set of primes dividing p together with the primes where the Tamagawa number of the Jacobian is divisible by p . Arguing as in the proof we see that, for primes v in $S \setminus S_1$, the local image $\Phi(\mathrm{Pic}_v^1(C))$ is a coset of the unramified subgroup of \mathcal{H}_v^0 . This observation can be used to improve the efficiency of the algorithm presented in section 3.

2. The local image

In order to use the description above to compute the algebraic Selmer set we need to be able to compute the image of the map

$$C(k_v) = \mathrm{Pic}_v^1(C) \xrightarrow{\Phi_v} \mathcal{H}_v$$

for a prime $v \in S$. Our primary interest is when p is an odd prime, the situation for $p = 2$ having received adequate attention elsewhere. For this reason we will ignore the archimedean primes. Even for odd p , the methods here should be familiar from p -descent on elliptic curves (or other objects). Perhaps the only substantial difference is that we work with \mathcal{H}_v rather than its image in $F_v^\times/k_v^\times F_v^{\times p}$.

If $(D, \pi) \in \mathrm{Cov}^{(p)}(C/k_v)$ is a p -covering of C defined over k_v with $D(k_v) \neq \emptyset$, then the image in $C(k_v)$ of the set of k_v -rational points on D is an orbit under the action of $p \mathrm{Pic}_v^0(C) = pE(k_v)$. On the other hand, Φ_v has the property that it takes a constant value on $\pi(D(k_v))$. Thus, Φ_v factors as

$$\Phi_v : \mathrm{Pic}_v^1(C) \rightarrow \mathrm{Pic}_v^1(C)/p \mathrm{Pic}_v^0(C) \hookrightarrow \mathcal{H}_v.$$

The second map is injective because the descent map is injective. Recall also that \mathcal{H}_v^0 sits in an exact sequence of \mathbb{F}_p -vector spaces

$$1 \rightarrow \frac{(\partial_2(\mu_p(\bar{F})))^{G_v}}{\partial_2((\mu_p(\bar{F})/\mu_p)^{G_v})} \rightarrow \mathcal{H}_v^0 \rightarrow \frac{F_v^\times}{k_v^\times F_v^{\times p}},$$

that \mathcal{H}_v is a coset and that the space on the left is finite.

All of this is actually valid over any perfect field of characteristic not equal to p . Over k_v , the sets $\text{Pic}_v^1(C)/p \text{Pic}_v^0(C)$ and \mathcal{H}_v are finite as well. As in the case of p -descent on elliptic curves, we can a priori determine the size of the former, and hence the size of the local image.

LEMMA 2.1. *For a non-archimedean prime v of k , $\# \frac{\text{Pic}_v^1(C)}{p \text{Pic}_v^0(C)} = p^d \cdot \#E(k_v)[p]$, where d is either $[k_v : \mathbb{Q}_p]$ or 0, correspondingly as v does or does not lie over p .*

PROOF: Since $C(k_v) \neq \emptyset$, $\frac{\text{Pic}_v^1(C)}{p \text{Pic}_v^0(C)}$ is a coset of $\frac{\text{Pic}_v^0(C)}{p \text{Pic}_v^0(C)}$, the size of which was given by lemma I.4.5. \square

To compute the local image it will thus suffice to find the images of sufficiently many independent points. It will actually be easier to determine independence by considering the images in \mathcal{H}_v . This is valid since the descent map is injective. Moreover, since the descent map is affine it suffices to find a set of images in \mathcal{H}_v which span a(n affine) space of the appropriate dimension. In practice we simply compute the images of random points until their images generate a large enough space. Using the fact that $C(k_v)$ is locally compact, one could also develop a deterministic algorithm for doing this (see for example, [St1] where the analogous situation is considered for 2-descent on Jacobians of hyperelliptic curves).

When working with objects defined over k_v (e.g. a point $P \in C(k_v)$, an element of H_v , etc.) in practice, we of course mean that we work up to some finite precision. To find a random point, we first choose a random solution to the defining equations of C modulo (some small power of) v . We then attempt to (randomly) lift to a solution modulo some higher power of v , applying Hensel's lemma along the way to guarantee that we are on the right track. This works quite well in practice. One can very quickly compute points up to very high precision.

Thus we can reduce the problem of computing the local image to the problem of deciding whether two points $P, Q \in C(k_v)$ (represented up to arbitrary precision) have the same image in \mathcal{H}_v . Evaluating $\tilde{\ell}$ on a tuple of v -integral elements in k representing homogeneous coordinates for a point $P \in C(k_v)$ up to some precision yields an element of H^\times which is v -adically close to some element of $\tilde{\mathcal{H}}_v$ representing $\Phi_v(P) \in \mathcal{H}_v$. That this is sufficient for our purposes follows from Hensel's lemma. The prototypical result in this direction is the following lemma.

LEMMA 2.2. *Suppose that q and r are rational primes and that $v|q$ with ramification index e . If $v \mid r$, then let n be the greatest integer $\leq e/(q-1) + e + 1$. Otherwise let $n = 1$. Then a v -adic unit $\alpha \in k_v^\times$ is an r -th power if and only if it is a r -th power modulo v^n .*

PROOF: This is [DSS, Lemma 13]. Note that this follows directly from Hensel's lemma when $v \nmid r$. When $v \mid r$, Hensel's lemma would lead to a bound $n = 2e + 1$. \square

COROLLARY 2.3. $F_v^\times/F_v^{\times p}$ is a finite-dimensional \mathbb{F}_p -vector space. The restriction map $F^\times \rightarrow F_v^\times/F_v^{\times p}$ is surjective and any two elements of F^\times which are sufficiently v -adically close to one another will have the same image.

PROOF: This is an obvious consequence once one notes that F_v splits as a product of completions of extensions of k at primes above v . Surjectivity then follows from the Chinese remainder theorem. \square

The image of k_v^\times in $F_v^\times/F_v^{\times p}$ is a subspace (whose dimension can be computed using the lemma). In practice, we ‘represent’ $F_v^\times/k_v^\times F_v^{\times p}$ as an \mathbb{F}_p -vector space W_1 and a homomorphism $\psi_1 : F^\times \rightarrow W_1$ giving the restriction map. In other words, the kernel of ψ_1 is the group of elements in F^\times which are of the form ad^p for some $a \in k_v^\times$, $d \in F_v^\times$. It remains to deal with the second component of H_v .

The kernel of the projection of $H_v^\times/k_v^\times \partial F_v^\times$ onto $F_v^\times/k_v^\times F_v^{\times p}$ is the finite \mathbb{F}_p -vector space

$$\mathcal{K}_v = \frac{(\partial_2(\mu_p(\bar{F})))^{G_v}}{\partial_2((\mu_p(\bar{F})/\mu_p)^{G_v})} \subset \frac{\mu_p(H_{2,v})}{\partial_2((\mu_p(\bar{F})/\mu_p)^{G_v})}.$$

To deal with this, we write down a map $\psi_2 : H_2 \rightarrow W_2$ to an \mathbb{F}_p -vector space W_2 that is isomorphic to $\mu_p(H_{2,v})/\partial_2((\mu_p(\bar{F})/\mu_p)^{G_v})$, with the following property. If $\eta \in H_2$ is sufficiently v -adically close to a p -th root of unity $\tilde{\eta} \in \mu_p(H_{2,v})$, then $\psi_2(\eta)$ is the corresponding element of W_2 .

Together the maps ψ_1 and ψ_2 allow us to test whether an element $(\delta, \varepsilon) \in \tilde{\mathcal{H}}_v^0$, represented up to sufficiently high precision by an element of H^\times , is trivial modulo $k_v^\times \partial F_v^\times$. First we check if $\psi_1(\delta) = 0 \in W_1$. If not, then (δ, ε) is nontrivial. If yes, then there exist $a \in k_v^\times$, $d \in F_v^\times$ (which we can compute up to sufficiently high precision) such that $\delta = ad^p$. Then

$$\varepsilon^p = \partial_2(\delta) = \partial_2(ad^p) = a^p \partial_2(d)^p.$$

It follows that $\eta := a\partial_2(d)/\varepsilon$ is a p -th root of unity. Then (δ, ε) is trivial modulo $k_v^\times \partial F_v^\times$ if and only if $\psi_2(\eta) = 0 \in W_2$.

To determine whether or not $P, Q \in C(k_v)$ have the same image under Φ_v , we check if $\tilde{\ell}(P)/\tilde{\ell}(Q)$ is trivial modulo $k_v^\times \partial F_v^\times$. Since $\tilde{\ell}(P)/\tilde{\ell}(Q) \in \tilde{\mathcal{H}}_v^0$, this can be done using ψ_1 and ψ_2 as above. In practice, we store the image $\tilde{\ell}(P_0) \in H^\times$ of a single point, together with ψ_1 , ψ_2 and subspaces of W_1 , W_2 from which we can recover the local image. This allows us to manipulate the local image using linear algebra. At any step in the process it is a simple matter to determine the space spanned by the images found up to that point. In particular, having computed the local image we can easily check whether a given element of H^\times restricts into $\Phi_v(\text{Pic}_v^1(C))$.

3. Algorithms

The theory above gives rise to the following algorithm for computing a set of representatives for the algebraic p -Selmer set of C . The output is a collection of elements in H^\times . Using the methods of section II.7, these can then be turned quite easily into explicit models as genus one normal curves of degree p^2 . Thus we have an algorithm for performing explicit second p -descents. For $p = 3$ and $k = \mathbb{Q}$, this has been implemented by the author in MAGMA and appears to perform quite well.

For larger p this is currently impractical for two reasons. The first of these is the, largely unavoidable, computation of S -class and -unit group information in F . Generically, the running time here is (at least) exponential in p^2 (see the discussion in section I.2). With the current state of the art, this becomes somewhat prohibitive already for $p = 5$. There is hope, however, that this will become feasible for larger p in the near future as computing power and algorithms in algebraic number theory improve.

The second arises from the fact that the algebra H_2 is simply too large. Generically it is a number field of degree $p^2(p^2 - 1)/2$ over k . The algorithm does not, however, require class and unit group information in H_2 . The most expensive operations required are the extraction of p -th roots. Even so, this quickly becomes impractical.

Compute $\text{Sel}_{alg}^{(p)}(C/k)$.

- (1) Compute the algebras F and H_2 , the map $\partial_2 : F \rightarrow H_2$, the linear form $\tilde{\ell}$, the constants $c \in k^\times$, $\beta \in H_2^\times$, and the set S of bad primes.
- (2) Let $V_1 \subset F^\times$ be a (finite) set of representatives for the unramified outside S subgroup of $F^\times/k^\times F^{\times p}$.
- (3) Let $V_2 = \{\delta \in V_1 : N_{F/k}(\delta) \equiv c \pmod{\mathbb{Q}^{\times p}}\}$.
- (4) For each $v \in S$, determine the local image $\Phi(\text{Pic}_v^1(C)) \subset \mathcal{H}_v$.
- (5) Let $V_3 = \{\delta \in V_2 : \forall v \in S, \text{res}_v(\delta) \in \text{pr}_1(\Phi(\text{Pic}_v^1(C)))\}$.
- (6) Let V_4 be the set of $(\delta, \varepsilon) \in F^\times \times H_2^\times$ such that $\delta \in V_3$ and $\varepsilon \in H_2^\times$ is a p -th root of $\partial_2(\delta)/\beta$, modulo the equivalence

$$(\delta, \varepsilon) \sim (\delta, \varepsilon') \Leftrightarrow \varepsilon/\varepsilon' \in \partial_2((\mu_p(\bar{F})/\mu_p)^{G_k}).$$
- (7) Let $V_5 = \{(\delta, \varepsilon) \in V_4 : \forall v \in S, \text{res}_v(\delta, \varepsilon) \in \Phi(\text{Pic}_v^1(C))\}$.
- (8) return V_5 .

REMARK: The reason for including steps (3) and (5) is to reduce the size of V_1 as much as possible before proceeding to step (6) where one has to extract p -th roots. Note that V_3 contains (but is not necessarily equal to) $\text{Sel}_{fake}^{(p)}(C/k)$. Using 1.5 we see that if $\#V_i \cdot \#\mathcal{K}_k < \#\text{Sel}^{(p)}(E/k)$ for some $i \in \{1, 2, 3\}$, then $\text{Sel}^{(p)}(C/k) = \emptyset$. In the typical applications $\#\mathcal{K}_k = 1$ and $\#\text{Sel}^{(p)}(E/k) \geq p^2$. In this way one can often confirm that $\text{Sel}^{(p)}(C/k) = \emptyset$ without running the full algorithm (cf. examples 2 and 3 in section 4).

Let us prove that the algorithm returns a set of representatives for the algebraic Selmer set. The equivalence in step (6) is included to ensure that the $(\delta, \varepsilon) \in V_5$ represent distinct classes modulo $k^\times \partial F^\times$. In the proof of 1.7 we have seen that, for primes not in S , a class in $H^\times/k^\times \partial F^\times$ is unramified if and only if its image in $F^\times/k^\times F^{\times p}$ is unramified. Moreover the elements of V_5 are in $\tilde{\mathcal{H}}_k$ by lemma II.5.2, since they restrict to $\tilde{\mathcal{H}}_v$ for some $v \in S \neq \emptyset$. It follows that V_5 is a set of representatives for $\{\Delta \in \mathcal{H}_S : \text{res}_v(\Delta) \in \text{Pic}_v^1(C), \forall v \in S\}$, which is equal to $\text{Sel}_{alg}^{(p)}(C/k)$ by Proposition

1.8.

REMARK: As remarked at the end of section 1 it may be possible to get away with a smaller set of bad primes. Prior to step (2) one can perform the following check. For each $v \in S$ that does not lie above p and for which the Tamagawa number of the Jacobian is not divisible by p , take any point on $C(k_v)$ and compute its image under Φ_v . Check if the image is unramified. If yes, then as the earlier remark indicated, the local image at v is equal to the unramified subgroup. So v can be removed from the set of bad primes.

Let us describe each step in more detail.

Step 1. It is clear what needs to be done in principle. We give the details for $p = 3$ assuming that the action of G_k on X is generic (by which we mean that the representation $G_k \rightarrow \text{AGL}(X)$ giving the action is surjective). There are finitely many other cases (for $p = 3$) which can all be handled similarly.

We store the algebras as quotients of the polynomial ring $k[\tau]$, so to compute them means to find a defining polynomial. Under the assumption that X is generic, F and H_2 are number fields of degrees 9 and 12 over k , respectively. We find a defining polynomial $f(\tau)$ for F by looking for the common zeros of U and h_U , where U is the cubic form defining C and h_U is the determinant of the Hessian matrix of U . This also gives a generic flex point x with coordinates in F . We choose a linear form $\tilde{\ell}_1 \in F[u_1, u_2, u_3]$ defining the tangent line to C at x and scale appropriately to make its coefficients integral (and reduce the number of common divisors if possible).

To compute H_2 , we adjoin a second root of $f(\tau)$ to F . This gives a field M_1 of degree 72 over k which contains the coordinates of a second flex x' . We then write down a linear form $\tilde{\ell}_2$ with coefficients in M_1 defining the line through x and x' . We choose some affine patch of \mathbb{P}^2 containing both x and x' in which the slope $m \in M_1$ of $\tilde{\ell}_2$ is well-defined. Since the action on X is generic, the twelve conjugates of this line all have distinct slopes. It follows that H_2 can be defined by the minimal polynomial of m over k .

For whatever reason, it tends to be much faster (in MAGMA) to first compute the minimal polynomial of m over F . This defines an extension M of F of relative degree 4. It is the étale algebra corresponding to the G_k -set of the 36 pairs² $(x, y) \in X \times Y_2$ consisting of a flex point x lying on a line y . We then compute the minimal polynomial of $m \in M$ over k , which gives a polynomial defining H_2 . Adjoining a root of $f(\tau)$ to H_2 gives M as an extension of H_2 . We may now coerce the coefficients of $\tilde{\ell}_2$ into H_2 .

The computation of these minimal polynomials is somewhat expensive, but it has the advantage that we obtain ∂_2 along the way. It is the composition $F \hookrightarrow M \rightarrow H_2$, where the second map is the norm of M/H_2 . The constant $\beta \in H_2^\times$ is given by the image of $\partial_2(\tilde{\ell}_1)/\tilde{\ell}_2^3$ in the coordinate ring $H_2[C]$ (i.e. modulo the ideal in $H_2[u_1, u_2, u_3]$ generated by $U = U(u_1, u_2, u_3)$). Similarly, the constant c can be computed as the image of $N_{F/k}(\tilde{\ell})/h_U^3$ modulo the ideal in $k[u_1, u_2, u_3]$ generated by U .

²Note that this algebra also plays a role in obtaining defining equations for the corresponding p -covering. See section II.7

The set of bad primes is the set of primes dividing c , the discriminant³ of U , or the discriminant of F , together with the primes above which there is a prime of F which divides all coefficients of $\tilde{\ell}_1$.

Step 2. This step is the main bottle-neck in the computation. Let \mathcal{F}_S denote the unramified outside S subgroup of $F^\times/k^\times F^{\times p}$. It fits into an exact sequence

$$k(S, p) \rightarrow F(S, p) \rightarrow \mathcal{F}_S \rightarrow \frac{\text{Cl}(\mathcal{O}_{k,S})}{p \text{Cl}(\mathcal{O}_{k,S})} \rightarrow \frac{\text{Cl}(\mathcal{O}_{F,S})}{p \text{Cl}(\mathcal{O}_{F,S})}.$$

For a derivation of this sequence and a description of how to compute \mathcal{F}_S see [PS, 12.8]. To compute it in practice we use the function `pSelmerGroup()` in MAGMA. Using this we compute a finite-dimensional \mathbb{F}_p -vector space $V_1 \simeq \mathcal{F}_S$ together with a set of representatives in F^\times .

Step 3. Since we have already computed $k(S, p)$, this can be accomplished very quickly using linear algebra over \mathbb{F}_p .

Step 4. Computing the Local Images. This was described in the previous section.

Step 5. The restriction map $\mathcal{F}_S \rightarrow F_v^\times/k_v^\times F_v^{\times p}$ is linear and given explicitly by ψ_1 (defined in the previous section), so this can be accomplished using linear algebra of \mathbb{F}_p .

Step 6. Extracting the p -th roots is straightforward (if a bit costly - it is here that the degree of H_2 becomes a problem). By ‘modulo the equivalence...’ we mean that we keep one (δ, ε) in each equivalence class. To determine equivalence, one needs to determine $(\mu_p(\bar{F})/\mu_p)^{G_k}$ and its image under ∂_2 . When $p = 3$, we are fortunate in that $\partial_2(\mu_p(F)) = \partial_2((\mu_p(\bar{F})/\mu_p)^{G_k})$ for all but one of the 46 possible Galois types⁴ (i.e. conjugacy classes of subgroups of $\text{AGL}_2(\mathbb{F}_3)$). So usually one need only determine $\mu_p(F)$ and its image under ∂_2 , which is easy. In the one exceptional case, the map $\mathcal{H}_k \rightarrow F^\times/k^\times F^{\times p}$ is injective. So after all is said and done one can simply identify all pairs of the form (δ, ε) and (δ, ε') . It is also worth mentioning that this exceptional case does not occur unless k contains the cube roots of unity. In particular it can be ignored when working over \mathbb{Q} .

Step 7. This is accomplished as in step (5) and as described in the previous section.

4. Examples

Full second 3-descent. We consider the elliptic curve

$$E : y^2 = x^3 + 3844x - 238328$$

labelled 61504bq1 in Cremona’s database. The torsion subgroup of $E(\mathbb{Q})$ is trivial, one quickly finds the point $(93, 961) \in E(\mathbb{Q})$ of infinite order so the rank is at least 1. A 2-descent shows that the rank is at most 1. This is in accordance with the analytic information coming from the L -series. The conjecture of BSD predicts that the order of $\text{III}(E/\mathbb{Q})$ is 9. The action of $G_{\mathbb{Q}}$ on $E[3]$ is given by the full general linear group

³This is the discriminant of the model as a genus one normal curve. It is somewhat unclear how one would go about computing this for an arbitrary prime p .

⁴This is also useful when writing down the map ψ_2 used to determine the local images.

$\mathrm{GL}_2(\mathbb{F}_3)$. In particular, descent by 3-isogeny is not possible. A full 3-descent on E produces a list of $13 = \frac{3^3-1}{2}$ plane cubics each representing an inverse pair of non-trivial elements in $\mathrm{Sel}^{(3)}(E/\mathbb{Q}) \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$,

$$\begin{aligned}
C_1 : 0 &= u_1^3 + 2u_2^3 + 2u_3^3 + 3u_1^2u_2 - 8u_1^2u_3 + 2u_1u_2^2 + 2u_1u_2u_3 + 23u_1u_3^2 + 27u_2u_3^2, \\
C_2 : 0 &= u_1^3 + 2u_2^3 + 10u_3^3 + 11u_1^2u_2 - 6u_1^2u_3 + 4u_1u_2^2 - 8u_1u_2u_3 - 10u_1u_3^2 - 2u_2^2u_3 + 4u_2u_3^2, \\
C_3 : 0 &= u_1^3 + 2u_2^3 + 30u_3^3 + 4u_1^2u_2 - 5u_1^2u_3 + 4u_1u_2u_3 + 2u_1u_3^2 - 2u_2^2u_3 + 12u_2u_3^2, \\
C_4 : 0 &= u_1^3 + 2u_2^3 + 59u_3^3 + 2u_1^2u_2 + u_1^2u_3 - 3u_1u_2^2 + 5u_1u_3^2 + 5u_2^2u_3 - 4u_2u_3^2, \\
C_5 : 0 &= u_1^3 + 3u_2^3 + 19u_3^3 - 3u_1^2u_2 + 2u_1^2u_3 + 7u_1u_2^2 + 4u_1u_2u_3 - 5u_1u_3^2 - 7u_2u_3^2, \\
C_6 : 0 &= u_1^3 + 6u_2^3 + 14u_3^3 + u_1^2u_2 + 5u_1^2u_3 - 6u_1u_2^2 + 4u_1u_2u_3 - 8u_1u_3^2 + 2u_2^2u_3, \\
C_7 : 0 &= u_1^3 + 7u_2^3 + 7u_3^3 - 3u_1^2u_2 + 7u_1^2u_3 + 4u_1u_2^2 + 4u_1u_2u_3 - 6u_1u_3^2 + 7u_2^2u_3 + 3u_2u_3^2, \\
C_8 : 0 &= 2u_1^3 + 0u_2^3 + 9u_3^3 + 8u_1^2u_2 + 4u_1^2u_3 + 8u_1u_2^2 - 4u_1u_2u_3 - 2u_1u_3^2 + 7u_2^2u_3 + 2u_2u_3^2, \\
C_9 : 0 &= 2u_1^3 + 2u_2^3 + 9u_3^3 + 2u_1^2u_3 + 10u_1u_2^2 - 4u_1u_2u_3 - u_1u_3^2 - 2u_2^2u_3 + 8u_2u_3^2, \\
C_{10} : 0 &= 2u_1^3 + 2u_2^3 + 13u_3^3 + 2u_1^2u_3 + 2u_1u_2^2 + 8u_1u_2u_3 + 19u_1u_3^2 - 8u_2^2u_3 + 6u_2u_3^2, \\
C_{11} : 0 &= 2u_1^3 + 4u_2^3 + 6u_3^3 - 4u_1^2u_2 + 14u_1^2u_3 - u_1u_2^2 - 8u_1u_3^2 - 3u_2^2u_3 + 2u_2u_3^2, \\
C_{12} : 0 &= 2u_1^3 + 4u_2^3 + 6u_3^3 + 6u_1^2u_2 + 12u_1^2u_3 - 9u_1u_2^2 + 2u_1u_3^2 + 5u_2^2u_3 + 4u_2u_3^2, \\
C_{13} : 0 &= 3u_1^3 + 3u_2^3 + 13u_3^3 - 4u_1^2u_2 + 2u_1^2u_3 + 5u_1u_2^2 + 2u_1u_2u_3 - 7u_1u_3^2 - 3u_2^2u_3 + 3u_2u_3^2.
\end{aligned}$$

Since we know the rank, the 3-descent implies that $\mathrm{III}(E/\mathbb{Q})[3] \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$. Moreover, exactly one of these cubics has a \mathbb{Q} -rational point. It is not hard to spot the point $(0 : 1 : 0)$ on the curve C_8 (the coefficient 0 is not a typographical error). The other 12 curves are counter-examples to the Hasse principle (each class in $(\mathrm{III}(E/\mathbb{Q})[3] \setminus \{0\})/\{\pm 1\}$ is represented by exactly 3 of these). We can verify this with a second 3-descent. This will also show that the 3-primary part of $\mathrm{III}(E/\mathbb{Q})$ is isomorphic to $\mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ as predicted by BSD. Note that it is enough to show that $\mathrm{Sel}^{(3)}(C_i/\mathbb{Q}) = \emptyset$ for a single i (see lemma I.5.7).

The curve $C := C_{10}$ has the flex algebra of smallest discriminant, so we start there. A defining polynomial for F is

$$f(t) = t^9 - 15t^8 + 12t^7 - 144t^6 - 12t^5 - 72t^4 + 390t^3 - 6t^2 - 570t + 354,$$

and the discriminant is $2^{16} \cdot 3^9 \cdot 29^6 \cdot 31^{14}$ (the prime 29 does not actually ramify in F). The flex points are given by

$$(1/1798(-12\tau^8 + 179\tau^7 - 119\tau^6 + 1575\tau^5 + 301\tau^4 - 854\tau^3 - 6174\tau^2 - 3478\tau + 6412) : \tau : 1),$$

where τ is a root of $f(t)$. The tangent line to C at this flex is defined by the linear form

$$\begin{aligned}
\tilde{\ell}_1 &= (5565\tau^8 - 77078\tau^7 - 24148\tau^6 - 799062\tau^5 - 943144\tau^4 \\
&\quad - 1168972\tau^3 + 1426282\tau^2 + 1938104\tau - 1717796)u_1 \\
&\quad + (50\tau^8 - 766\tau^7 + 782\tau^6 - 6614\tau^5 \\
&\quad\quad + 2232\tau^4 + 6404\tau^3 + 26402\tau^2 + 29152\tau - 39538)u_2 \\
&\quad + (8351\tau^8 - 123761\tau^7 + 72967\tau^6 - 1121941\tau^5 - 275125\tau^4 \\
&\quad\quad + 237690\tau^3 + 4395686\tau^2 + 2301490\tau - 5326066)u_3,
\end{aligned}$$

which is integral and has good reduction⁵ outside $\{2, 29, 31\}$. The constant $c \in \mathbb{Q}^\times$ is determined by taking the ratio of $N_{F/\mathbb{Q}}(\tilde{\ell}_1)$ and the cube of the Hessian of C and considering its image modulo the ideal generated by the cubic defining C . We find that c is congruent to $2^2 \cdot 31$ modulo cubes. So the set of ‘bad primes’ is the set $S = \{2, 3, 29, 31\}$.

Using \mathcal{F}_S to denote the unramified outside S part of $F^\times/\mathbb{Q}^\times F^{\times 3}$ it follows that the image of the 3-Selmer set of C under the fake descent map is contained in the affine space

$$V_{c,S} = \{\delta \in \mathcal{F}_S : N_{F/\mathbb{Q}}(\delta) \equiv 124 \pmod{\mathbb{Q}^{\times 3}}\}$$

(this is the set labelled V_2 in step (3) of the algorithm in section 3). Since the action of $G_{\mathbb{Q}}$ on the flex points is generic, we know that the projection $\mathcal{H}_{\mathbb{Q}} \rightarrow F^\times/\mathbb{Q}^\times F^{\times 3}$ is injective. So, in fact, \mathcal{H}_S is mapped injectively to the set

$$V_{\beta,S} = \{\delta \in \mathcal{F}_S : \partial_2(\delta) \equiv \beta \pmod{H_2^{\times 3}}\} \subset V_{c,S}.$$

Computing $V_{c,S}$ using MAGMA’s `pSelmerGroup()` function takes about a minute. We find that it is an affine space of dimension 6 over \mathbb{F}_3 .

One now has options for moving forward. The algorithm outlined in the previous section would proceed to compute local images at the primes of S . The alternative would be to compute the set $V_{\beta,S}$. This would require either computing $H_2(S, 3)$ and doing some linear algebra or checking whether or not lots of elements in H_2^\times are cubes. The latter is feasible and leads to the conclusion⁶ that $\mathcal{H}_S = \emptyset$. This means that there are no 3-coverings of C (with trivial obstruction) which are locally solvable everywhere outside S . In particular, the 3-Selmer set is empty.

On the other hand, computing the local images is much faster since one can avoid working too much with H_2 . By considering the decomposition of the primes in the 3-division field of E and using lemma 2.1 we determine that

$$\#\Phi_v(C(\mathbb{Q}_v)) = \begin{cases} 1, & \text{for } v \in \{2, 31\}, \\ 3, & \text{for } v \in \{3, 29\}. \end{cases}$$

The defining polynomial of F has a linear factor over \mathbb{Q}_2 , so there is a \mathbb{Q}_2 -rational flex point on C . Using lemma II.6.2 we conclude that the fake descent map is injective over \mathbb{Q}_2 . We compute that $F_2^\times/\mathbb{Q}_2^\times F_2^{\times 3}$ is a 1-dimensional \mathbb{F}_3 -vector space and that the restriction map $\text{res}_2 : V_{c,S} \rightarrow F_2^\times/\mathbb{Q}_2^\times F_2^{\times 3}$ is constant. We easily find a \mathbb{Q}_2 -point on C whose image under $\Phi_{fake,2}$ is the same, so the prime 2 gives us no information. Geometrically, this says that any⁷ 3-covering of C which is locally solvable everywhere outside S is also locally solvable at 2. Since there is a \mathbb{Q}_{31} -rational flex, the local information at 31 can be collected similarly. In this case however, the subspace of $V_{c,S}$ restricting into the image of $C(\mathbb{Q}_{31})$ has codimension 2.

For $v \in \{3, 29\}$, the local images have size 3 and the projection pr_1 in the following commutative diagram is no longer injective.

$$\begin{array}{ccc} C(\mathbb{Q}_v) & & \\ \Phi_v \downarrow & \searrow \Phi_{fake,v} & \\ \mathcal{H}_v & \xrightarrow{\text{pr}_1} & F_v^\times/\mathbb{Q}_v^\times F_v^{\times 3} \end{array}$$

⁵By scaling it should actually be possible to obtain a linear form with good reduction at 29 as well.

⁶One can show, however, that $\mathcal{H}_{\mathbb{Q}} \neq \emptyset$.

⁷Since we know $\mathcal{H}_S = \emptyset$ this condition is vacuous.

For $v = 29$, the kernel of the projection has dimension 1 and the image of $C(\mathbb{Q}_{29})$ in \mathcal{H}_{29} consists of a full fiber over a point in $F_{29}^\times/\mathbb{Q}_{29}^\times F_{29}^{\times 3}$. On the other hand, the image of $C(\mathbb{Q}_3)$ in \mathcal{H}_3 is spread across three fibers (i.e. the image in the lower-right space is of size 3). In principle, one might have to now consider the fibers above elements of $V_{c,S}$ under the projection $\text{pr}_1 : \mathcal{H}_{\mathbb{Q}} \rightarrow F^\times/\mathbb{Q}^\times F^\times$ (this would mean computing \mathcal{H}_S , which would involve taking cube roots) and check their images in \mathcal{H}_3 . In this case however, the images of the \mathbb{Q}_v -points in $F_v^\times/\mathbb{Q}_v^\times F_v^{\times 3}$ impose enough local conditions to show that

$$\text{Sel}_{\text{fake}}^{(3)}(C/\mathbb{Q}) \subset \{\delta \in V_{c,S} : \forall v \in S, \text{res}_v(\delta) \in \Phi_{\text{fake},v}(C(\mathbb{Q}_v))\} = \emptyset,$$

which shows that $\text{Sel}^{(3)}(C/\mathbb{Q})$ is also empty.

In this example, the algorithm of section 3 terminates after step 5 without ever having to extract cube roots. This somewhat fortuitous situation seems to occur quite often (at least when the Selmer set in question is actually empty), but even when it does one may still need to make use of H_2 . We needed it to compute the local image at 29. By way of a different example, we consider the curve C_1 . Here we find (somewhat surprisingly) that the constant $c \in \mathbb{Q}^\times$ (coming from the condition on $N_{F/\mathbb{Q}}(\tilde{\ell}_1)$) is actually a cube⁸. The action on X is again generic, so the fake descent map is globally injective. This means we have bijections (cf. II.5.6)

$$\begin{aligned} \text{Cov}_0^{(3)}(C_1/\mathbb{Q}) &\simeq \{\delta \in F^\times/\mathbb{Q}^\times F^{\times 3} : \partial_2(\delta)/\beta \in H_2^{\times 3}\}, \text{ and} \\ C_1^\perp &\simeq \{\delta \in F^\times/\mathbb{Q}^\times F^{\times 3} : \partial_2(\delta) \in H_2^{\times 3}\} \end{aligned}$$

We can compute H_2 and β , and check that $\beta \notin H_2^{\times 3}$ (so the sets above are distinct). However, since $c \in \mathbb{Q}^{\times 3}$ both sets are contained in

$$\{\delta \in F^\times/\mathbb{Q}^\times F^{\times 3} : N_{F/\mathbb{Q}}(\delta) \in \mathbb{Q}^{\times 3}\}.$$

This shows that the algebra H_2 is needed to describe the image even when the fake descent map is injective.

Producing elements of order 9 in III. Consider the curve

$$E : y^2 + xy = x^3 - 1479474x - 692765778$$

labelled 5514a3 in Cremona's database. His tables indicate that the analytic rank of E is 0 and that the analytic order of the Shafarevich-Tate group is 81. The torsion subgroup is trivial and a 2-descent confirms that $E(\mathbb{Q}) = \{0_E\}$. A 3-descent produces the 4 plane cubic curves

$$\begin{aligned} C_1 : x^3 + 6y^3 + 919z^3 &= 53xyz, \\ C_2 : 2x^3 + 3y^3 + 919z^3 &= 53xyz, \\ C_3 : x^3 + 3y^3 + 1838z^3 &= 53xyz, \\ C_4 : x^3 + 2y^3 + 2757z^3 &= 53xyz. \end{aligned}$$

These evidently correspond to the four ways of choosing an unordered triple $\{a, b, c\}$ of distinct⁹ positive integers such that $abc = 2 \cdot 3 \cdot 919$. Each C_i represents an inverse

⁸The analogous situation occurs for $p = 2$ if and only if the double cover of \mathbb{P}^1 has a pair of rational points above ∞ . However, the curve C_1 has no rational points.

⁹The unique unordered triple of nondistinct positive integers with this property corresponds to the curve $C : x^3 + y^3 + 5514z^3 = 55xyz$ which represents the trivial element in $\text{Sel}^{(3)}(E/\mathbb{Q})$. It has the obvious (and unique) \mathbb{Q} -point $(1 : -1 : 0)$.

pair of nontrivial elements in $\text{Sel}^{(3)}(E/\mathbb{Q})$, so

$$\text{Sel}^{(3)}(E/\mathbb{Q}) \simeq \text{III}(E/\mathbb{Q})[3] \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$$

and each C_i is a counter-example to the Hasse principle.

Since the analytic order of $\text{III}(E/\mathbb{Q})$ is 81, we expect that $\#\text{Sel}^{(3)}(C_i/\mathbb{Q}) = 9$ for $i = 1, \dots, 4$. This is confirmed by performing a second 3-descent. Note that using lemma I.5.7 it suffices to do the computation for a single i . Each of the 9 elements computed in $\text{Sel}_{alg}^{(3)}(\mathbb{Q}, C_i)$ correspond to a pair of inverse elements of order 9 in $\text{Sel}^{(9)}(E/\mathbb{Q}) \simeq \text{III}(E/\mathbb{Q})[9] \simeq \mathbb{Z}/9\mathbb{Z} \times \mathbb{Z}/9\mathbb{Z}$. Using the methods of section II.7 we can explicitly compute models for these coverings as genus one normal curves of degree 9 in \mathbb{P}^8 defined by the vanishing of 27 quadrics. To our knowledge this is the largest prime power¹⁰ to date for which such examples have been produced. An example of an everywhere locally solvable 3-covering of C_1 is given on the opposite page. Our algorithm originally produced models with absurdly large coefficients. With the help of Tom Fisher and Michael Stoll we have managed to get the coefficients down to a reasonable size. While there is little room for improvement in this particular example, it would be nice to develop better techniques of minimization and reduction for these models (cf. section 5).

¹⁰Presumably Fisher can produce examples of order 12 by combining examples of orders 3 and 4.

Defining equations for a curve in \mathbb{P}^8 representing an element of $\text{Sel}^{(3)}(C_1/\mathbb{Q})$

$$\begin{aligned}
0 &= z_2 z_5 - z_5 z_6 + 3z_7^2 + z_7 z_9 + 2z_8^2, \\
0 &= z_1 z_2 + z_2 z_6 + z_2 z_7 + z_4 z_9 + z_5 z_7 + 2z_5 z_8, \\
0 &= -z_1 z_7 + z_2 z_8 + z_4 z_5 + z_5^2 + 2z_6 z_7 + z_6 z_8 - z_7^2, \\
0 &= z_2^2 + z_2 z_6 - 2z_4 z_7 - z_4 z_8 - 2z_5 z_7 + 2z_5 z_8 - z_5 z_9, \\
0 &= -z_1 z_7 - z_2 z_7 + z_2 z_8 - z_5^2 - 2z_6 z_7 - z_6 z_8 - z_6 z_9 - z_7^2, \\
0 &= z_1 z_8 + 3z_2 z_7 - z_2 z_8 + z_3 z_5 + z_4 z_5 + z_5^2 + z_5 z_7 + z_7 z_8, \\
0 &= -z_1 z_6 + z_2 z_6 + z_3 z_7 + 2z_4 z_7 - z_4 z_8 - 2z_5 z_8 - z_6 z_7 + z_7^2, \\
0 &= -z_2 z_5 + z_3 z_6 + z_5 z_6 + z_6 z_7 + z_7^2 + 2z_7 z_8 + 2z_7 z_9 - z_8 z_9, \\
0 &= -z_1 z_5 + z_3 z_6 + z_5 z_6 - z_5 z_7 + z_6 z_7 - 2z_7^2 - z_7 z_8 + 2z_7 z_9 + 2z_8^2, \\
0 &= -z_1 z_4 - 2z_1 z_5 + z_2 z_3 + z_2 z_4 - z_2 z_5 + z_2 z_7 - z_4 z_7 + 3z_5 z_6 - 2z_5 z_7, \\
0 &= -z_1 z_7 - z_1 z_8 + z_2 z_7 - z_2 z_8 - z_2 z_9 + z_3 z_5 - z_5^2 + z_5 z_7 - z_7^2 - z_7 z_8, \\
0 &= z_1 z_2 + z_1 z_6 - z_2 z_6 + z_2 z_7 - z_4 z_7 + z_4 z_9 - 2z_5 z_7 + z_5 z_9 - z_6^2 + z_6 z_7, \\
0 &= -z_1 z_2 + z_2^2 - z_2 z_7 + z_3 z_7 + 2z_3 z_8 + z_4 z_7 + z_4 z_8 + z_5 z_7 + z_7^2 + 2z_7 z_8, \\
0 &= 3z_1 z_5 - z_3 z_6 + 3z_4 z_6 + 2z_5 z_6 + 3z_5 z_7 - z_6 z_7 - z_7^2 + 3z_7 z_8 - 2z_7 z_9 + z_8^2, \\
0 &= z_1 z_9 + z_2 z_8 - 4z_2 z_9 + z_3^2 + z_3 z_4 + 2z_3 z_7 + z_4 z_7 - z_5^2 - z_6 z_8 + 2z_6 z_9 + z_7^2 + z_7 z_9, \\
0 &= -z_1 z_6 + z_2^2 - z_3 z_7 - z_3 z_8 - z_3 z_9 + 2z_4 z_7 - z_4 z_9 - z_5 z_7 - 2z_6^2 - z_6 z_7 - z_7^2 - z_7 z_8 - z_7 z_9, \\
0 &= -2z_1 z_8 - z_1 z_9 + 2z_2 z_7 - z_2 z_8 + z_2 z_9 - z_3 z_5 + z_4^2 + z_5^2 - z_5 z_7 - z_6 z_8 + 2z_6 z_9 - 2z_7 z_8 - z_7 z_9, \\
0 &= z_1 z_2 + z_1 z_6 + z_2 z_7 - z_3 z_7 + z_3 z_8 + z_4 z_7 + z_4 z_8 + z_4 z_9 - z_5 z_7 - 2z_5 z_8 - 2z_5 z_9 + z_6 z_7 - z_7^2 + z_7 z_8, \\
0 &= z_1 z_5 + z_2 z_3 + z_2 z_4 - z_2 z_5 + z_2 z_7 + z_3 z_6 + z_4 z_6 + z_5 z_7 + z_6 z_7 + z_7^2 - z_7 z_8 + z_7 z_9 - 2z_8^2 + 2z_8 z_9 - z_9^2, \\
0 &= -z_1 z_8 + 2z_1 z_9 - 2z_2 z_7 - z_2 z_9 + z_3 z_4 + z_3 z_5 + z_4 z_7 + z_5^2 + z_5 z_7 + z_6 z_7 - z_6 z_8 - 2z_6 z_9 - z_7 z_8 + 2z_7 z_9, \\
0 &= -z_1^2 + z_1 z_2 + 2z_1 z_6 - z_1 z_7 - 2z_2^2 + z_2 z_6 + z_2 z_7 - z_2 z_8 - z_4 z_5 - z_4 z_7 - z_5^2 + 2z_5 z_7 - z_5 z_8 - z_6^2 - z_6 z_8, \\
0 &= -z_1 z_7 + z_1 z_8 - z_1 z_9 - 3z_2 z_7 - 2z_2 z_8 - z_3 z_4 - z_3 z_5 + z_4 z_5 - z_4 z_7 - z_5 z_7 - z_6 z_7 + z_6 z_8 - z_7^2 + z_7 z_8 - z_7 z_9, \\
0 &= -z_1 z_3 - z_1 z_4 - z_1 z_5 - z_1 z_7 - 2z_2 z_4 - 3z_2 z_5 - z_3 z_7 - z_4 z_7 + z_5 z_6 - z_5 z_7 + z_7^2 - z_7 z_8 - z_7 z_9 - z_8^2 - 2z_8 z_9 + z_9^2, \\
0 &= -z_1 z_8 + z_1 z_9 + z_2 z_8 + z_2 z_9 + z_3 z_4 + z_3 z_5 + z_4^2 + 2z_4 z_5 + z_4 z_7 + z_5^2 + z_5 z_7 - 2z_6 z_7 + 2z_6 z_8 - 2z_6 z_9 - z_7 z_8 + z_7 z_9, \\
0 &= z_1 z_3 - z_1 z_5 + z_2 z_3 + 2z_2 z_4 + z_2 z_8 + z_3 z_6 + z_3 z_7 - z_4 z_6 - z_5^2 - z_5 z_7 - z_6 z_7 - z_6 z_8 - z_6 z_9 + z_7^2 - z_7 z_8 - z_7 z_9 - z_8^2 \\
&\quad - 2z_8 z_9 - 2z_9^2, \\
0 &= z_1^2 - 2z_1 z_2 - z_1 z_6 - z_1 z_8 + z_2 z_6 - z_2 z_7 - z_2 z_9 + z_3 z_5 + 2z_3 z_7 + z_3 z_9 + z_4 z_5 + z_4 z_8 - z_5 z_7 + 2z_5 z_8 + z_6^2 \\
&\quad + z_6 z_7 + z_6 z_8 + z_7^2 - z_7 z_8 + z_7 z_9, \\
0 &= z_1 z_2 - z_1 z_8 - z_2^2 + z_2 z_3 - 2z_2 z_5 - z_2 z_7 + z_2 z_8 - z_3 z_5 - z_3 z_7 - 2z_3 z_8 - z_4 z_5 - z_4 z_7 - z_4 z_8 - z_5^2 - 2z_5 z_7 + z_7^2 \\
&\quad + 2z_7 z_8 - 2z_7 z_9 + 2z_8^2 + 2z_8 z_9 - z_9^2.
\end{aligned}$$

An example of second 5-descent. Consider the genus one normal curve of degree 5

$$C = \left\{ \begin{array}{l} u_1^2 - 7u_3u_4 + 2u_2u_5 = 0 \\ u_2^2 - 21u_4u_5 + u_1u_3 = 0 \\ u_3^2 - 3u_1u_5 + u_2u_4 = 0 \\ 7u_4^2 - u_1u_2 + 2u_3u_5 = 0 \\ 6u_5^2 - u_2u_3 + u_1u_4 = 0 \end{array} \right\} \subset \mathbb{P}^4.$$

It represents an inverse pair of elements in $\text{III}(E/\mathbb{Q})[5]$, where E is the elliptic curve,

$$E : y^2 + xy = x^3 - 109388x - 13934358.$$

labelled 1050o2 in Cremona's database. This elliptic curve is somewhat special in that it admits a rational 5-isogeny. Such curves were studied by Fisher in his PhD thesis [Fi2]; we have taken C from his website. His computations show that $\text{III}(E/\mathbb{Q})[5] = \text{Sel}^{(5)}(E/\mathbb{Q})$ and that both are of dimension 2 over \mathbb{F}_5 . We perform a 5-descent on C to show that $\text{Sel}^{(5)}(C/\mathbb{Q}) = \emptyset$ and consequently that $\text{III}(E/\mathbb{Q})[5^\infty] = \text{III}(E/\mathbb{Q})[5]$, i.e. that there is no higher-powered 5-torsion in $\text{III}(E/\mathbb{Q})$.

The special $G_{\mathbb{Q}}$ -structure of the set of flex points C makes the computations practical. The flex points of C split into 5 $G_{\mathbb{Q}}$ -orbits, each consisting of 5 flexes. Each orbit is the intersection of one of the coordinate hyperplanes, $\{u_i = 0\}$, with C . For example, the hyperplane $\{u_5 = 0\}$ intersects C in the points

$$(-\theta^3 : -\theta^2 : \theta : 1 : 0)$$

where θ is a 5-th root of 7. The corresponding factor of F is $\mathbb{Q}(\theta)$.

The splitting of the flex algebra also leads to a splitting of H_2 , the largest factors being of degree 25. While there is a considerable amount of 'book keeping' involved in keeping track of all the factors, computations in H_2 should be entirely practical. This would be necessary if one were to construct an explicit model of some 5-covering of C , but to show that $\text{Sel}^{(5)}(C/\mathbb{Q}) = \emptyset$ it will be enough to work only with the fake descent map.

LEMMA 4.1. *The projection $\mathcal{H}_{\mathbb{Q}} \rightarrow F^\times/\mathbb{Q}^\times F^{\times 5}$ is injective.*

PROOF: Of the 132 possible Galois actions on the flex points, the map fails to be injective for only 28 (cf. the table at the end of section II.6). A direct computation shows that the action on X is not of one of these 28 types. \square

REMARK: For at least one of the 28 types for which injectivity fails, the action of Galois factors through a cyclic subgroup of the Galois group of the splitting field of X . Chebotarëv's density theorem then implies that injectivity fails to hold over \mathbb{Q}_v for infinitely many primes v .

From the lemma it follows that the 5-Selmer set of C maps injectively to the fake 5-Selmer set. Since the 5-Selmer group of the Jacobian has dimension 2, we know that $\text{Sel}^{(5)}(C/\mathbb{Q})$ is either of dimension 2 or is empty (see prop. 1.4 and cor. 1.5). To show that $\text{Sel}^{(5)}(C/\mathbb{Q}) = \emptyset$ it suffices to show that the image of $\text{Sel}^{(5)}(C/\mathbb{Q})$ under the fake descent map has fewer than 25 elements.

Write $F \simeq F_1 \times \cdots \times F_5$, each factor corresponding to the flex points on the hyperplane $\{u_i = 0\}$. Use N_i to denote the norm from F_i to \mathbb{Q} . The linear form $\tilde{\ell}_1 \in F[u_1, \dots, u_5]$ used to define the fake descent map splits as a tuple $\tilde{\ell}_1 = (L_1, \dots, L_5)$

of linear forms, each defined over some F_i . It is easy to see that $N_i(L_i)$ and u_i^5 cut out the same divisors on C . This leads to relations of the form $N_i(L_i) = c_i \cdot u_i^5$ in the coordinate ring of C , where $c_i \in \mathbb{Q}^\times$ are some constants. Arguing as in section II.5 we see that the image of $\text{Cov}_0^{(5)}(C/\mathbb{Q})$ under the fake descent map is contained in the set

$$V_c := \{(\delta_1, \dots, \delta_5) = \delta \in \frac{F^\times}{\mathbb{Q}^\times F^{\times 5}} : N_i(\delta_i) \equiv c_i \pmod{\mathbb{Q}^{\times 5}}\}.$$

REMARK: This is slightly more restrictive than II.5 where we showed that the image of the fake descent map is contained in the set $\{\delta \in F^\times/\mathbb{Q}^\times F^{\times 5} : N(\delta) \equiv \prod_i c_i \pmod{\mathbb{Q}^{\times 5}}\}$.

For each $i \in \{1, \dots, 5\}$, we can easily compute such an L_i . For example, the hyperplane meeting $(-\theta^3 : -\theta^2 : \theta : 1 : 0)$ with multiplicity 5 is defined by the vanishing of

$$L_5 := 8\theta u_1 + 5\theta^2 u_2 + 5\theta^3 u_3 + 56\theta^4 u_4 + 41u_5.$$

From this one can compute that $c_5 \equiv 7 \pmod{\mathbb{Q}^{\times 5}}$. Similar computations for the other i show that the only primes appearing in the factorization of some c_i with multiplicity prime to 5 are contained in the set $S = \{2, 3, 5, 7\}$. This set also contains all primes of bad reduction for C and the L_i . Hence, the fake Selmer set is unramified outside S .

Use \mathcal{F}_S to denote the unramified outside S part of $F^\times/\mathbb{Q}^\times F^{\times 5}$. From the discussion above it follows that the image of $\text{Sel}^{(5)}(C/\mathbb{Q})$ under the fake descent map is contained in the set

$$V_{c,S} := \{(\delta_1, \dots, \delta_5) = \delta \in \mathcal{F}_S : N_i(\delta_i) \equiv c_i \pmod{\mathbb{Q}^{\times 5}}\}.$$

Since F splits, computation of the unramified outside S subgroup is fast. The subset $V_{c,S}$ is a translate¹¹ of the space

$$V_{1,S} := \{(\delta_1, \dots, \delta_5) = \delta \in \mathcal{F}_S : N_i(\delta_i) \equiv 1 \pmod{\mathbb{Q}^{\times 5}}\}.$$

This turns out to be a 2-dimensional \mathbb{F}_5 -vector space, generated by the classes of the elements $d_1, d_2 \in F^\times \simeq F_1^\times \times \dots \times F_5^\times$ given by

$$\begin{aligned} d_1 &:= (5, 5, 5, 5, 19 + 12\theta - 4\theta^2 - 7\theta^3 - \theta^4), \text{ and} \\ d_2 &:= (5, 5, 5, 5, 414744 + 314727\theta - 58629\theta^2 - 169062\theta^3 - 43881\theta^4), \end{aligned}$$

where θ is a 5-th root of 7.

In order to cut down the size any further we need to use information coming from some prime in S . Considering the decomposition of the prime 2 in the constituent fields of F one can see that there is exactly one flex point defined over \mathbb{Q}_2 . Existence of a \mathbb{Q}_2 -rational flex point means that $E[5]$ and X are isomorphic as $G_{\mathbb{Q}_2}$ -sets. Using lemma 2.1 this tells us that the local image, $\Phi_{fake,2}(C(\mathbb{Q}_2)) \subset F_2^\times/\mathbb{Q}_2^\times F_2^{\times 5}$, consists of a single element. In order that an element in $V_{c,S}$ lie in the fake Selmer set it is necessary that it restrict to this element. One can check, however, that the restrictions of d_1 and d_2 above are nontrivial. This means that the image of $V_{c,S}$ in $F_2^\times/\mathbb{Q}_2^\times F_2^{\times 5}$ is of size at least 5. It follows that the image of the 5-Selmer set under $\tilde{\Phi}_{fake}$ is a proper subset of $V_{c,S}$. As such it has strictly fewer than 25 elements, proving that $\text{Sel}^{(5)}(C/\mathbb{Q}) = \emptyset$.

¹¹One can check that $V_{c,S} \neq \emptyset$

Verifying BSD. As a final example we offer the following theorem which, in addition to some very deep results regarding BSD, brings to bear many of the currently available algorithms for explicit descents on curves of genus one.

THEOREM 4.2. *The full Birch and Swinnerton-Dyer conjecture holds for the elliptic curves*

$$\begin{aligned} E : y^2 &= x^3 + 7^3 \cdot 61^3 \cdot 97^4, \text{ and} \\ E' : y^2 &= x^3 - 3^3 \cdot 7^3 \cdot 61^3 \cdot 97^4 \end{aligned}$$

defined over \mathbb{Q} . In particular, $\text{III}(E/\mathbb{Q}) \simeq \text{III}(E'/\mathbb{Q}) \simeq \mathbb{Z}/12\mathbb{Z} \times \mathbb{Z}/12\mathbb{Z}$.

REMARK: Along the lines of theorem I.1.1 it should be possible to explicitly write down a list of $74 = \frac{12^2-4}{2} + 4$ genus one normal curves in projective space with the property that if V/\mathbb{Q} is any variety such that for all $p \leq \infty$, $V \otimes \mathbb{Q}_p$ is \mathbb{Q}_p -isomorphic to $E \otimes \mathbb{Q}_p$, then V is \mathbb{Q} -isomorphic to exactly one curve in the list. The models for the elements of order 2, 3 and 4 are produced by the descents described below. We would then need Fisher's method [Fi5] for combining these to produce models for the elements of order 6 and 12.

As in Theorem I.1.1, the hard part of the proof is taken care of by the existing partial results in the direction of BSD. The role of descent is to compute the p -primary parts of the Shafarevich-Tate groups at the primes 2 and 3. Note that these curves are related by the 3-isogeny

$$h : E \ni (a, b) \mapsto \left(\frac{a^3 + 2^2 7^3 61^3 97^4}{a^2}, \frac{a^3 b - 2^3 7^3 61^3 97^4 b}{a^3} \right) \in E'.$$

So the validity of BSD for either curve implies its validity for the other.

One can check that the values of the L -series of E and E' at $s = 1$ are (equal and) approximately 5.5542. Results of Coates and Wiles then imply that the Mordell-Weil groups are finite [CoWi]. One easily checks that there is no nontrivial torsion on either, so the Mordell-Weil groups are trivial. The predicted orders of $\text{III}(E/\mathbb{Q})$ and $\text{III}(E'/\mathbb{Q})$ are the numbers

$$\begin{aligned} \text{III}_{an}(E) &= \frac{L_E(1)}{\Omega(E) \cdot \prod_{p|\Delta(E)} C_p(E)} \text{ and} \\ \text{III}_{an}(E') &= \frac{L_{E'}(1)}{\Omega(E') \cdot \prod_{p|\Delta(E')} C_p(E')}, \end{aligned}$$

where $L(s)$ is the L -series, Ω is the real period, C_p denotes the Tamagawa number at p and Δ is the discriminant (note that the regulators and torsion subgroups are trivial).

The real period of E is $\Omega(E) \approx 0.0096427$ and the only Tamagawa number not equal to 1 is $C_7(E) = 4$. This gives

$$\text{III}_{an}(E) \approx \frac{5.5542}{(0.0096427) \cdot 4} \approx 144$$

The real period of E' satisfies $\Omega(E) = 3 \cdot \Omega(E')$ and the nontrivial Tamagawa numbers are $C_3(E') = 3$ and $C_7(E') = 4$. Thus $\text{III}_{an}(E') \approx 144$ as well.

It is known that $\text{III}_{an}(E)$ and $\text{III}_{an}(E')$ are rational numbers of (explicitly) bounded denominator, so (taking the computations to sufficiently high precision) we conclude that they are in fact equal to 144. Rubin's result then implies that $\text{III}(E/\mathbb{Q})[p] = 0$ for

all primes p not dividing $\text{III}_{an}(E)$ or the order of the group of units in the ring of integers of the field of complex multiplication. The same holds for $\text{III}(E'/\mathbb{Q})$. These curves have CM by $\sqrt{-3}$, so we conclude that both Shafarevich-Tate groups are annihilated by some power of 6.

After applying these deep results, we are left only with the task of computing the 2- and 3-primary parts of $\text{III}(E/\mathbb{Q})$ and $\text{III}(E'/\mathbb{Q})$. Since the validity of BSD for a given elliptic curve is actually a property of its isogeny class, it will suffice to perform the computations for either curve. We describe the computations for E . The computations for E' are similar (and equally feasible).

The 2-primary part. One needs explicit first and second 2-descents to produce models for the elements of order dividing 4 and then a third 2-descent to show that there are no elements of order 8. The 2-descent on E yields models for the 3 nontrivial elements of $\text{III}(E/\mathbb{Q})[2]$ as double covers of \mathbb{P}^1 :

$$\begin{aligned} C_1 : u_3^2 &= 130174u_1^4 - 71004u_1^3 - 426024u_1^2 + 2011780u_1 - 390522, \\ C_2 : u_3^2 &= 11834u_1^4 + 260348u_1^3 - 710040u_1^2 + 1372744u_1 + 3999892, \\ C_3 : u_3^2 &= 5917u_1^4 + 29585u_1^3 - 177510u_1^2 + 804712u_1 + 562115. \end{aligned}$$

For each C_i a second 2-descent will produce a pair of quadric intersections, each representing a pair of inverse elements of order 4 in $\text{III}(E/\mathbb{Q})$. For example, $\text{Sel}^{(2)}(C_3/\mathbb{Q})$ is of order 4 and represented by the two curves

$$\begin{aligned} D_1 &= \left\{ \begin{array}{l} 2z_1^2 + 14z_1z_2 - 3z_2^2 + 4z_1z_3 - 2z_2z_3 + 5z_3^2 + 8z_1z_4 + 2z_2z_4 - 8z_3z_4 - 15z_4^2 = 0 \\ 24z_1^2 + 8z_1z_2 - 22z_2^2 + 36z_1z_3 + 18z_2z_3 + 63z_3^2 - 54z_1z_4 - 24z_2z_4 + 42z_3z_4 + 14z_4^2 = 0 \end{array} \right\} \subset \mathbb{P}^3, \\ D_2 &= \left\{ \begin{array}{l} 3z_1^2 + 2z_1z_2 + 3z_2^2 + 6z_1z_3 - 2z_2z_3 - 8z_3^2 + 6z_1z_4 + 24z_2z_4 - 13z_4^2 = 0 \\ 6z_1^2 + 86z_1z_2 - 20z_2^2 - 18z_1z_3 + 2z_2z_3 + 13z_3^2 - 18z_1z_4 - 22z_2z_4 - 6z_3z_4 - 42z_4^2 = 0 \end{array} \right\} \subset \mathbb{P}^3. \end{aligned}$$

One then uses Stamminger's method for third 2-descent which shows that none of the elements of order 4 lift to elements of order 8. It follows that the 2-primary part is $\text{III}(E/\mathbb{Q})[2^\infty] \simeq \mathbb{Z}/4\mathbb{Z} \times \mathbb{Z}/4\mathbb{Z}$.

The 3-primary part. For this we can make use of the 3-isogeny. A 3-isogeny descent computes that $\text{Sel}^{(h)}(E/\mathbb{Q}) \simeq \text{Sel}^{(h)}(E'/\mathbb{Q}) \simeq \mathbb{Z}/3\mathbb{Z}$. Since the Mordell-Weil groups are trivial these Selmer groups are isomorphic to the corresponding torsion subgroups of the Shafarevich-Tate groups. The exact sequence

$$0 \rightarrow \frac{E'(\mathbb{Q})[h']}{h(E(\mathbb{Q})[3])} \rightarrow \text{Sel}^{(h)}(E/\mathbb{Q}) \rightarrow \text{Sel}^{(3)}(E/\mathbb{Q}) \xrightarrow{h} \text{Sel}^{(h')}(E'/\mathbb{Q}) \rightarrow \frac{\text{III}(E'/\mathbb{Q})[h']}{h(\text{III}(E/\mathbb{Q})[p])} \rightarrow 0$$

reduces to

$$0 \rightarrow \text{Sel}^{(h)}(E/\mathbb{Q}) \rightarrow \text{Sel}^{(3)}(E/\mathbb{Q}) \rightarrow \text{Sel}^{(h')}(E'/\mathbb{Q}) \rightarrow 0,$$

which splits since $\text{Sel}^{(3)}(E/\mathbb{Q})$ is 3-torsion. We conclude that $\text{III}(E/\mathbb{Q})[3] \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$.

The 3-isogeny descent (implemented in MAGMA) is also explicit in that it produces the projective plane cubic

$$C : u_1^3 + 4u_2^3 + 4017643u_3^3 = 4u_1^2u_2 + 3u_1u_2^2$$

representing the pair of nontrivial elements in $\text{Sel}^{(h)}(E/\mathbb{Q})$. This also represents an inverse pair of nontrivial elements in $\text{Sel}^{(3)}(E/\mathbb{Q})$. In order to show that $\text{III}(E/\mathbb{Q})[3^\infty] \simeq \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ it will be enough to show that $\text{Sel}^{(3)}(C/\mathbb{Q}) = \emptyset$.

For this we do a second 3-descent. The reducibility of $E[3]$ translates into a splitting of the flex algebra. We find that F is isomorphic to the product of the cubic and sextic number fields with defining polynomials

$$f(t) = t^3 - 4t^2 - 3t + 4 \text{ and}$$

$$f(t) = t^6 + 2408704t^3 + 5533080062500$$

and the set of bad primes is $S = \{2, 3, 5, 7, 61, 97\}$. Despite the splitting, computation of $F(S, 3)$ takes a couple hours of processor time. Having accomplished that however the remaining computations are very fast. The fake 3-Selmer set is contained in the set of all $\delta \in F^\times / \mathbb{Q}^\times F^{\times 3}$ such that

- (1) δ is unramified outside S ,
- (2) $N_{F/\mathbb{Q}}(\delta) \equiv 7^2 61^2 97 \pmod{\mathbb{Q}^{\times 3}}$ and
- (3) $\forall v \in S, \text{res}_v(\delta) \in \Phi_{\text{fake},v}(C(\mathbb{Q}_v))$.

This set can be computed without ever using H_2 and turns out to be empty. It follows that the 3-Selmer set of C is empty and thus that the 3-primary part of $\text{III}(E/\mathbb{Q})$ is isomorphic to a product of two cyclic groups of order 3.

5. Directions for further work

Minimization and reduction. As we have noted, our algorithm currently produces models that may be utterly useless for finding rational points. The problem being that we have not developed a systematic way of ensuring that the coefficients of these models are of a reasonable size. While there is much room for both theoretical and practical improvement, we have had some success in a few preliminary examples. The primary source of inspiration is the paper of Cremona, Fisher and Stoll [CFS] where the problem of obtaining nice models for genus one normal curves of degrees 2, 3 and 4 is solved.

The process breaks into two steps termed *minimization* and *reduction*. The first works locally at a non-archimedean prime and aims to remove prime factors from the invariants of the model. In practice we have had some success using the ad hoc methods developed by Fisher when he was working with 6- and 12-coverings. But in order to develop the theory in any comprehensive way, one would first need to develop a theory of invariants for genus one normal curves of larger degree. At the moment this seems difficult, partly because it is not entirely clear what the appropriate definition of an integral genus one model of degree n should be. Naively one could define this as a collection of $n(n-3)/2$ quadrics in n variables with integral coefficients, but an arbitrarily chosen collection has a poor chance of defining a genus one normal curve. In any event, the invariants will be given as polynomials in the coefficients of the model and should be related to the standard invariants, c_4, c_6 and the discriminant, of the Jacobian. These polynomials will likely be too large to ever write down, so one may also need a clever way of evaluating them.

In contrast, the outlook for reduction is somewhat more optimistic. After having removed prime factors from the invariants, one wants to make a unimodular transformation to make the coefficients smaller (it is somewhat enlightening to think of this as minimization at an archimedean prime). Here there is a clear theoretical description over \mathbb{Q} for arbitrary n , which reduces the problem to lattice reduction. The question is how to go about doing it in practice. Fisher has explained to me how the method of

reduction for 4-coverings generalizes to 9-coverings, though at the time of writing this had not been implemented.

Even higher descents. Once one develops a reasonable method for producing nice models of p^2 -coverings of elliptic curves one might try to do descent on these. Stamminger’s method of third 2-descent could be used as a guide for a general method for third p -descents. One step in his method involves solving a conic over a cubic number field. The analog for larger p will likely require finding a rational point on a $p - 1$ dimensional Brauer-Severi variety defined over a number field of degree $p^2 - 1$. Already for $p = 3$ this seems quite difficult. Even so a theoretical description would be nice. Stamminger’s method actually computes a fake 2-Selmer set. As with second descents, the naive generalization to odd p will likely omit too much information for one to be able to recover the genuine p -Selmer set. So there will likely be a bit of work involved in correcting for this.

In a different direction, one might try to combine the results of Stamminger’s third 2-descents with those of a second 3-descent. Fisher has devised a method for producing explicit models for 6- and 12-coverings of an elliptic curve starting from a pair of 2- and 3- or 3- and 4-coverings [Fi15]. He has used this to find generators of the Mordell-Weil group of unprecedentedly large height. It would seem that aspects of his method apply to any pair of consecutive prime powers, so one might try combining 8- and 9-coverings to produce 72-coverings of elliptic curves. For this to be useful (or even feasible) in practice one would, of course, first need to deal with minimization and reduction of 8- and 9-coverings.

Descent on higher genus curves. In retrospect, the criterion for choosing the family of functions defining our descent map was that it could be used to perform a p -descent on the Jacobian (cf. II.4.5). Namely, the induced map on $\text{Pic}_K^0(C)$ could be identified with the connecting homomorphism in the Kummer sequence of multiplication by p on the Jacobian. For genus one curves it is likely to be easier to perform descent on the Jacobian by using functions on the Jacobian itself. The point, however, is that any family of functions on a curve C of arbitrary genus which are useful for computing or bounding a Selmer group of the Jacobian should also allow one to obtain information on coverings of C . This is implicit in the approach taken by Bruin and Stoll [BS] who have used the $u_1 - \theta$ map to perform (fake) 2-descent on hyperelliptic curves, rather than on their Jacobians.

In the more general situation of cyclic covers of the projective line of the form $u_3^p = f(u_1)$, Poonen and Schaefer [PS] use the $u_1 - \theta$ map to perform fake descents on the Jacobian. Stoll and Van Luijk have suggested [StVL] that one can use additional information coming from the function $u_3 \in \kappa(C)^\times$ to ‘unfake the fake Selmer group’. They construct a homomorphism from the subgroup of $\text{Pic}(C)$ consisting of divisors whose degrees are divisible by p which eliminates the ambiguity in the fake descent. It would seem the techniques of this thesis can be used to induce a similar homomorphism on the full Picard group. This would then allow one to compute the corresponding Selmer set of C or to study coverings of the torsor in $H^1(K, \text{Jac}(C))$ corresponding to $\text{Pic}^1(C)$.

Period-index questions. In addition to being able to compute a potential covering obstruction to the existence of rational points, such explicit descriptions of the coverings could potentially be used to study subtle period-index questions for C . In the genus one case, the image of the descent map gives a parameterization of $\text{Cov}_0^{(p)}(C/K)$.

If this is nonempty, then C is divisible by p in the Weil-Châtelet group by an element of index dividing p^2 . The explicit description of \mathcal{H}_K allows one to relate existence to the solvability of certain norm equations (cf. II.5.4). We have been able to show, for an elliptic curve E over a number field k , that every element of $\text{III}(E/k)[2]$ is 2-divisible by an element in $H^1(k, E)$ of index dividing 4 (Theorem I.2.5). It would be interesting to determine whether this remains true for odd primes and, if not, then use the description to find a counterexample.

One could also study the analogous question for coverings of cyclic covers of the projective line. In this situation the coverings parameterized by the analogous descent map should be those for which the pull-back of a ramification point (which always defines a K -rational divisor class) is linearly equivalent to a K -rational divisor. Using the explicit description of the image of the (fake) descent map for hyperelliptic curves given in [BS], we have been able to construct examples of everywhere locally solvable hyperelliptic curves over \mathbb{Q} for which no such coverings exist [Cr2]. It seems that this somewhat unexpected result might be explained by the existence of everywhere locally solvable hyperelliptic curves over \mathbb{Q} which have no 2-coverings defined over \mathbb{Q} .

Bibliography

- [Ba] A. BANDINI: *Three-descent and the Birch and Swinnerton-Dyer conjecture*, Rocky Mountain J. Math. **34** (2004), 13-27.
- [BSD-I] B.J. BIRCH AND H.P.F. SWINNERTON-DYER, *Notes on elliptic curves. I.*, J. Reine Angew. Math. **212** (1963), 7-25.
- [BSD-II] —, *Notes on elliptic curves. II.*, J. Reine Angew. Math. **218** (1965), 79-108.
- [MAGMA] MAGMA, in W. BOSMA, J. CANNON AND C. PLAYOUST, *The Magma algebra system I: The user language*, J. Symbolic Comput. **24** (1997), 235-265, available online at <http://magma.maths.usyd.edu.au/magma>
- [BC] A. BREMNER AND J.W.S. CASSELS, *On the equation $Y^2 = X(X^2 + p)$* , Math. Comp. **42** (1984), 257-264.
- [BS] N. BRUIN AND MICHAEL STOLL, *Two-cover descent on hyperelliptic curves*, Math. Comp. **78** (2009), 2347-2370.
- [Ca1] J.W.S. CASSELS, *Arithmetic on curves of genus 1. II. A general result*, J. Reine Angew. Math. **203** (1960), 174-208.
- [Ca2] —, *Arithmetic on curves of genus 1. IV. Proof of the Hauptvermutung*, J. Reine Angew. Math. **211** (1962), 95-112.
- [Ca3] —: *Arithmetic on curves of genus 1. V. Two counter-examples*, J. Lond. Math. Soc. **38** (1963), 244-248.
- [Ca4] —, *Lectures on elliptic curves*, LMS Student Texts 24, Cambridge University Press, Cambridge, 1991.
- [Ca5] —, *Second descents for elliptic curves*, J. Reine Angew. Math. **494** (1998), 101-127.
- [ChWe] C. CHEVALLEY AND A. WEIL, *Un théorème d'arithmétique sur les courbes algébriques*, C. R. Acad. Sci. Paris **195** (1932), 570-572.
- [Cl] P.L. CLARK, *Period-index problems in WC-groups I: elliptic curves*, J. Number Theory **114** (2005), 193-208.
- [CS] P.L. CLARK AND S. SHARIF, *Period, index and potential Sha*, arXiv:math/0811.3019.
- [CoWi] J. COATES AND A. WILES: *On the conjecture of Birch and Swinnerton-Dyer*, Invent. Math. **39** (1977), 223-251.
- [CP] H. COHEN AND F. PAZUKI, *Elementary 3-descent by 3-isogeny*, Acta Arith. **140** (2009), 369-404.
- [CFS] J. CREMONA, T.A. FISHER AND M. STOLL, *Minimisation and reduction of 2- 3- and 4-coverings of elliptic curves*, arXiv:math/0908.1741.
- [CFOSS-I] J.E. CREMONA, T.A. FISHER, C. O'NEIL, D. SIMON AND M. STOLL, *Explicit n-descent on elliptic curves, I. Algebra*, J. Reine Angew. Math. **615** (2008), 121-155.
- [CFOSS-II] —, *Explicit n-descent on elliptic curves, II. Geometry*, J. Reine Angew. Math. **632** (2009), 63-84.
- [CFOSS-III] —, *Explicit n-descent on elliptic curves, III. Algorithms*, (in preparation).
- [Cr1] B. CREUTZ, *A Grunwald-Wang type theorem for abelian varieties*, (preprint).
- [Cr2] —, *Answer to a question of Bruin and Stoll*, (preprint).
- [DSS] Z. DJABRI, E.F. SCHAEFER AND N.P. SMART, *Computing the p-Selmer group of an elliptic curve*, Trans. Amer. Math. Soc. **352** (2000), 5583-5597.
- [Fi1] T.A. FISHER, *On 5 and 7 descents for elliptic curves*, PhD thesis, University of Cambridge, 2000.
- [Fi2] —, *Some examples of 5 and 7 descent for elliptic curves over \mathbb{Q}* , J. Eur. Math. Soc. **3** (2001), 169-201.

- [Fi3] —, *A counterexample to a conjecture of Selmer*, in: *Number theory and algebraic geometry*, M. Reid, A. Skorobogatov (eds.), LMS Lecture Note Series **303**, Cambridge University Press, Cambridge, 2003, 119-132.
- [Fi4] —, *Some improvements to 4-descent on an elliptic curve*, in: *Algorithmic number theory*, A.J. van der Poorten and A. Stein (eds.), Lecture Notes in Comput. Sci. **5011**, Springer, Berlin, 2008, 125-138.
- [Fi5] —, *Finding rational points on elliptic curves using 6-descent and 12-descent*, *J. Algebra* **320** (2008), 853-884.
- [Go] E.H. GOINS, *Explicit descent via 4-isogeny on an elliptic curve* (2004); arXiv:math/0411215v1
- [FG] E.V. FLYNN AND C. GRATTONI, *Descent via isogeny on elliptic curves with large rational torsion subgroups*, *J. Symbolic Comput.* **43** (2008), 293-303.
- [La] S. LANG, *Abelian varieties*, Springer, Berlin, 1983.
- [Le1] H. LENSTRA JR., *Galois theory for schemes*, course notes, available online at <http://www.websites.math.leidenuniv.nl/algebra>
- [Le2] —, *Algorithms in algebraic number theory*, *Bull. Amer. Math. Soc.* **26** (1992), 211-244.
- [Mat] A. MATTUCK, *Abelian varieties over p -adic ground fields*, *Ann. of Math.* **62** (1955), 92-119.
- [Maz] B. MAZUR, *On the passage from local to global in number theory*, *Bull. Amer. Math. Soc.* **29** (1993), 14-50.
- [Mor] L.J. MORDELL, *On the rational solutions of the indeterminate equations of the 3rd and 4th degrees*, *Proc. Camb. Phil. Soc.* **21** (1922), 179-192.
- [MSS] J.R. MERRIMAN, S. SIKSEK AND N.P. SMART, *Explicit 4-descents on an elliptic curve*, *Acta Arith.* **77** (1996), 385-404.
- [CoN] J. NEUKIRCH, A. SCHMIDT AND K. WINGBERG, *Cohomology of number fields* (SECOND EDITION), Grundlehren Math. Wiss. **323**, Springer, Berlin, 2008.
- [O'N] C. O'NEIL, *The period-index obstruction for elliptic curves*, *J. Number Theory* **95** (2002), 329-339.
- [PS] B. POONEN AND E.F. SCHAEFER, *Explicit descent for Jacobians of cyclic covers of the projective line*, *J. Reine Angew. Math.* **488** (1997), 141-188.
- [Ru] KARL RUBIN: *The "main conjectures" of Iwasawa theory for imaginary quadratic fields*, *Invent. Math.* **103** (1991), 25-68.
- [Ser1] J.P. SERRE, *Local fields*, Grad. Texts in Math. **67**, Springer, New York, 1979.
- [Ser2] —, *Algebraic groups and class fields*, Springer, New York, 1988.
- [Ser3] —, *Galois cohomology*, Springer, Berlin, 1997.
- [Sch1] E.F. SCHAEFER, *Class groups and Selmer groups*, *J. Number Theory* **56** (1996), 79-114.
- [Sch2] —, *Computing a Selmer group of a Jacobian using functions on the curve*, *Math. Ann.* **310** (1998), 447-471.
- [SchSt] E.F. SCHAEFER AND M. STOLL, *How to do a p -descent on an elliptic curve*, *Trans. Amer. Math. Soc.* **356** (2004), 1209-1231.
- [Sel] E.S. SELMER, *The diophantine equation $ax^3 + by^3 + cz^3 = 0$* . *Acta. Arith.* **85** (1951), 203-362.
- [Sik1] S. SIKSEK, *Descent on curves of genus 1*, PhD thesis, Exeter, 1995.
- [Sik2] —, *Infinite descent on elliptic curves*, *Rocky Mountain J. Math.* **25** (1995), 1501-1538.
- [Sik3] —, *Descent on Picard groups using function on curves*, *Bull. Aust. Math. Soc.* **66** (2002), 119-124.
- [Sil] J.H. SILVERMAN, *The arithmetic of elliptic curves*, Grad. Texts in Math. **106**, Springer, New York, 1986.
- [Sim] D. SIMON, *Computing the rank of elliptic curves over number fields*, *LMS J. Comput. Math.* **5** (2002), 7-17.
- [Sta] S. STAMMINGER, *Explicit 8-descent on elliptic curves*, PhD thesis, International University Bremen, 2005.
- [Ste] N.M. STEPHENS: *The diophantine equation $X^3 + Y^3 = DZ^3$ and the conjectures of Birch and Swinnerton-Dyer*, *J. Reine Angew. Math.* **231** (1968), 121-162.
- [St1] M. STOLL, *Implementing 2-descent for Jacobians of hyperelliptic curves*, *Acta Arith.* **98** (2001), 245-277.

- [St2] —, *Descent on elliptic curves*, course notes, available online at www.mathe2.uni-bayreuth.de/stoll
- [StVL] M. STOLL AND R. VAN LUIJK, *Unfaking the fake Selmer group*, (preprint).
- [Top] J. TOP, *Descent by 3-isogeny and 3-rank of quadratic fields*, in: *Advances in number theory*, F. Gouvea and N. Yui (eds.), Clarendon Press, Oxford, 1993, 303-317.
- [Weil] A. WEIL, *Sur un théorème de Mordell*, *Bull. Sci. Math.* **54** (1930), 182-191.
- [Wom] T. WOMACK, *Explicit descent on elliptic curves*, PhD thesis, Nottingham, 2003.