

An application of the p -adic analytic class number formula

Claus Fieker and Yinan Zhang

ABSTRACT

We propose an algorithm to verify the p -part of the class number for a number field K , provided K is totally real and an abelian extension of the rational field \mathbb{Q} , and p is any prime. On fields of degree 4 or higher, this algorithm has been shown heuristically to be faster than classical algorithms that compute the entire class number, with improvement increasing with larger field degrees.

1. Introduction

The quotient of the group of invertible ideals of a number field K , modulo principal ideals, is the class group of K , denoted by Cl_K . It is one of the fundamental invariants of the field, and of core importance to almost all multiplicative problems concerning number fields. As a result, the ability to compute Cl_K is an important task in algebraic number theory. Whilst there are conjectures about the structures of class groups, the computation of Cl_K is difficult and existing approaches to obtaining provable results are slow. These either assume some generalised Riemann hypothesis, thus delivering results that are not proven, or make use of Minkowski-type bounds, which is computationally infeasible for most examples.

There are, however, circumstances where only the p -part of the class group is required. This is especially important in certain areas of Iwasawa theory and elliptic curves, where they are used in descents to find rational points on elliptic curves. Here, it would be useful to have an algorithm that could efficiently compute only the p -part.

Whilst there have been approaches to this problem in the past, including attempts by Gras and Gras [8], much progress has been made in the past fifteen years, including most recently work by Hakkarainen [9], which focused on an algorithm to find prime divisors of class numbers, and Aoki and Fukuda [1], whose algorithm was more focused on p -adic decomposition of the class group. Both algorithms require the condition that p does not divide the field degree of K , and $p \neq 2$ (problematic as given a fixed degree, genus theory indicates that there are infinitely many fields with class number divisible by the degree). This prevents them from dealing with all fields K which are abelian extensions of the rational field \mathbb{Q} , despite a theoretical result from Leopoldt showing that this is possible [10, § 5.5].

In this paper we propose a new algorithm to compute the p -part of the class number for any totally real abelian number field K and prime p . The result is unconditional and can be used to verify the p -part of the class group. Just as classical algorithms use the class number formula for their computation, this algorithm makes use of the p -adic version of the formula. Whilst this may not be the most efficient way to implement a p -adic algorithm to compute the p -part of the class group, it does present an unconditional method that runs in polynomial time in the conductor of the field.

Received 2 October 2013; revised 7 January 2016.

2010 Mathematics Subject Classification 11S40 (primary), 11-04, 11Y35, 11Y70 (secondary).

The computation of the p -part of the class number, apart from few special cases, is usually done through a computation of the structure of the full class group using a variation of Buchmann’s subexponential algorithm. The method essentially proceeds in two steps. First, a (small) finite set of prime ideals is chosen. The algorithm then proceeds to determine the subgroup of the class group generated by those ideals. In the second step, the choice of the initial ideals is verified by checking all prime ideals of norm up to some bound.

Depending on the application, the bound can be of size $O(\log^2 |D|)$, where D is the discriminant of the number field, in the case where the generalised Riemann hypothesis (GRH) is assumed, or of size $O(\sqrt{|D|})$ for unconditional results. As a consequence the running time is overwhelmingly dominated by the verification step in all but the trivial examples. In this paper, we propose a new method that can compute unconditionally the p -part of the class number (under a reasonable heuristic assumption supported by numerical data) in time polynomial in $O(n^{-1}\sqrt{|D|})$ for cyclic fields of prime degree n . This allows an asymptotically much faster unconditional verification than any previously known method. At the end of the paper, we produce examples showing the approach to be practical as well.

2. p -adic class number formula

Our algorithm is based on the p -adic class number formula, which provides a link between the p -adic L -function, the p -adic regulator and class number of a number field [14, Theorem 5.24].

THEOREM 2.1. *Suppose K is a totally real abelian number field, with discriminant D , regulator R_p and class number h . Let its group of corresponding Dirichlet characters be X . Then*

$$\frac{2^{n-1}hR_p}{\sqrt{D}} = \prod_{\substack{\chi \in X \\ \chi \neq 1}} \left(1 - \frac{\chi(p)}{p}\right)^{-1} L_p(1, \chi), \tag{2.1}$$

where n is the field degree of K , up to choice of sign for \sqrt{D} .

Provided we are able to compute $L_p(1, \chi)$ for the required characters and R_p , it is possible for us to calculate h . To do so we make use of two formulae for computing $L_p(1, \chi)$. The first one [7, Theorem 11.5.37] is a closed formula in terms of (p -adic) logarithms, similar to the formula for $L(1, \chi)$:

THEOREM 2.2. *Let χ be an even character with conductor f_χ , and ζ a primitive f_χ th root of unity. If χ is the trivial character then $L_p(s, \chi)$ has a pole at $s = 1$. Otherwise*

$$L_p(1, \chi) = -\left(1 - \frac{\chi(p)}{p}\right) \frac{\sum_{a=1}^{f_\chi} \chi(a)\zeta^a}{f_\chi} \sum_{i=1}^{f_\chi} \bar{\chi}(i) \log_p(1 - \zeta^{-i}).$$

Note that $\sum_{a=1}^{f_\chi} \chi(a)\zeta^a$ is a Gauss sum.

The second [7, Proposition 11.3.8] is a convergent series:

THEOREM 2.3. *Let χ be a primitive character of conductor f_χ , let $m = \text{lcm}(f_\chi, q_p)$, where $q_p = 4$ if $p = 2$ and $q_p = p$ otherwise. If χ is a non-trivial character then $L_p(1, \chi)$ is given by the formula*

$$L_p(1, \chi) = \sum_{\substack{0 \leq a < m \\ (a, p) = 1}} \chi(a) \left(-\frac{\log_p(a)}{m} + \sum_{j \geq 1} (-1)^j \frac{m^{j-1} B_j}{a^j j} \right),$$

where B_j is the j th Bernoulli number.

The main steps in computing the p -adic L -functions involve computing p -adic logarithms and creation of the \mathbb{Q}_p extension fields required for parts of the formula. In addition, the appropriate characters for K have to be selected, and R_p calculated. We shall define these formally as we come to compute them.

3. Computing p -adic L -functions

Let p be a prime number. Denote by \mathbb{Q}_p the field of rational p -adic numbers, with the usual p -adic norm $|\cdot|_p$ and valuation v_p . Let $\overline{\mathbb{Q}_p}$ be the algebraic closure of \mathbb{Q}_p , and \mathbb{C}_p the topological closure of $\overline{\mathbb{Q}_p}$ with respect to $|\cdot|_p$.

3.1. p -adic logarithm

Evaluating $L_p(1, \chi)$ requires the use of the p -adic logarithm, \log_p , in both cases. It is defined by the usual power series expansion,

$$\log_p(1 + X) = \sum_{i=1}^{\infty} \frac{(-1)^{i+1} X^i}{i}.$$

Here the series has a radius of convergence of 1, so the domain of $\log_p(x)$ is $U_1 = \{a \in \mathbb{C}_p \mid |x - 1|_p < 1\}$.

It is possible to extend this logarithm to \mathbb{C}_p^\times . We know that by [14, Proposition 5.4], any element $x \in \mathbb{C}_p^\times$ can be uniquely represented in the form

$$x = p^r \omega u,$$

where r is some rational number, ω is a root of unity of order prime to p , and $u \in U_1$, and that there is a unique extension of \log_p from U_1 to \mathbb{C}_p^\times given by

$$\log_p(x) := \log_p(u).$$

REMARK 3.1. The above logarithm commutes with Frobenius endomorphism, which maps elements in a commutative ring of characteristic p to their p th powers.

A key problem in computing $L_p(1, \chi)$ is the need for the computation of p -adic logarithms of arbitrary elements. The straightforward power series is only valid for 1-units, that is, elements in $1 + p\mathbb{Z}_K$. Its naive use would require us to extend the field, which we want to avoid. As r is rational ($r = a/b$), the field needs to be extended to contain a uniformising element of valuation $(1/b)$.

ALGORITHM 3.2. *Computation of the p -adic logarithm of an arbitrary element x .*

Input: x

Output: $\log_p x$

- 1: $k := v_p(x)$ and $y := \pi^{-k}x$, where π is a uniformising element of $\mathbb{Q}_p(x)$
- 2: $z := y^{n-1}$ where $n := \#F$ for the residue class field F
- 3: Use the power series to compute $\log z$ and $\log y := 1/(n - 1) \log z$
- 4: $\varepsilon := \pi^e/p$
- 5: Return $\log_p x := (k/e) \log \varepsilon + \log y$.

Proof. We know that x can be rewritten as $p^r \omega u$. Let e be the ramification index of $\mathbb{Q}_p(x)/\mathbb{Q}_p$, and the valuation of x be v . Then we have $r = v/e$.

Let π be a uniformising element of $\mathbb{Q}_p(x)$, that is, an element with valuation 1. Then $\pi^e = p\varepsilon$, for some unit ε . Using this fact, we compute ε . Now we rewrite x so that

$$x = p^{v/e} \omega u \pi^{-v/e} = p^{v/e} u (p\varepsilon)^{-v/e} = u \varepsilon^{-v/e}.$$

Taking \log_p of both sides, we get

$$\log_p(x) = \log_p(u \varepsilon^{-v/e}) = \log_p(u) - \frac{v}{e} \log_p(\varepsilon).$$

Since $\log_p(x) = \log_p(u)$, we need to add a correction factor of $(v/e) \log_p(\varepsilon)$ to return the correct value, and this completes our algorithm. □

Recall that this logarithm commutes with the Frobenius endomorphism. This allows faster computations of the terms $\log_p(1 - \zeta^{-i})$ in Theorem 2.2 by making use of the Frobenius endomorphism (where applicable) to reduce the number of logarithms calculated, which is in general computationally tedious, especially in fields with large ramification degree. This maps $\log_p(1 - \zeta^i)$ to $\log_p(1 - \zeta^{ip})$ at a fraction of the cost of actually evaluating $\log_p(1 - \zeta^{ip})$.

We can also evaluate $1/p^l \log_p(1 + X)^{p^l}$ instead of $\log_p(1 + X)$. This reduces the number of terms that need to be computed from the power series, as it now converges more quickly due to larger valuations. However, this comes at the cost of requiring additional precision for the division by p^l . For each explicit example we can calculate the optimal value of l for the logarithm computation.

An example of this is when $p = 3$ and ζ is a 1423rd root of unity. The time taken to compute $\log_p(1 - \zeta^{-1200})$ is around 53 s without this optimisation (that is, $l = 0$), decreasing to 27 s when $l = 15$ before increasing again with larger values of l .

Bernstein outlined a fast algorithm for logarithms [2] which may be used here. This, along with other fast algorithms for logarithms, are intended for calculations requiring high precision (at least several thousand digits) and it is not clear whether this is required for our algorithm.

3.2. \mathbb{Q}_p extension field creation

We need to construct a field that enables us to compute $L_p(1, \chi)$. Both approaches require roots of unity, either explicitly in the calculation or for the construction of the Dirichlet character χ . This calls for a cyclotomic extension of \mathbb{Q}_p . Suppose we need an o th root of unity for the calculations. We can write

$$o = p^{v_p(o)} c$$

so that p and c are relatively prime. The value of c and valuation of o determine whether we need a ramified or unramified extension, or possibly both.

The following algorithm constructs the necessary extension.

ALGORITHM 3.3. Construction of an extension field of \mathbb{Q}_p containing a o th root of unity

Input: o, \mathbb{Q}_p

Output: $\mathbb{Q}_p[\zeta_o]$

- 1: $c := o \pmod{p^{v_p(o)}}$
- 2: $f := \text{order of } p \pmod{c} \text{ or } 1 \text{ if } c = 1$
- 3: Construct T , unramified extension of \mathbb{Q}_p of degree f
- 4: If $v_p(o) > 0$, $g(x) := ((x+1)^p - 1)/(x)$
- 5: Construct S , totally ramified extension of T defined by $g(x)$
- 6: If $v_p > 1$ then $h(x) := g(x+1)^{v_p(o)-1} - \pi - 1$, where π is a uniformiser
- 7: Construct R , totally ramified extension of S defined by $h(x)$.

This is achieved through the construction of the intermediate fields below:

$$\begin{array}{c}
 \mathbb{Q}_p[\zeta_o] = \mathbb{Q}_p[\zeta_{p^{v_p(o)}}, \zeta_c] \\
 \swarrow \quad \searrow \\
 \mathbb{Q}_p[\zeta_p, \zeta_c] \\
 \swarrow \quad \searrow \\
 \mathbb{Q}_p[\zeta_c] \\
 \downarrow \\
 \mathbb{Q}_p
 \end{array}$$

Proof. $\mathbb{Q}_p[\zeta_c]$ is an unramified extension of \mathbb{Q}_p . Since \mathbb{Q}_p contains the $(p-1)$ th roots of unity, an unramified extension of degree f would yield (p^f-1) th roots of unity. It is clear that if $n|(p^f-1)$, then the smallest such f is the order of p modulo c , by definition.

The construction of the totally ramified extensions is simple once we have $\mathbb{Q}_p[\zeta_c]$. The only thing we need to be careful about is to ensure that the defining polynomials are Eisenstein. Since $g(x)$ is the p th cyclotomic polynomial, evaluating it at $x+1$ instead of x satisfies our criterion. A similar process follows for $h(x)$. \square

3.3. Computing $L_p(1, \chi)$

The computation based on Theorem 2.2 is mostly clear. On the other hand, when using the formula from Theorem 2.3, it is important to know how many terms in the infinite sum need to be calculated for the result have precision ν , that is, correct in value modulo p^ν . This follows as a corollary from the proof that the infinite sum converges.

PROPOSITION 3.4. *The infinite sum*

$$\sum_{j \geq 1} (-1)^j \frac{m^{j-1} B_j}{a^j j}$$

converges with respect to $|\cdot|_p$.

Proof. Let s_j be the j th term of the sequence. Since $|\cdot|_p$ is a non-Archimedean norm it is sufficient to show that $\lim_{j \rightarrow \infty} s_j = 0$.

Consider the valuation of the individual terms in s_j . Since $(a, p) = 1$,

$$v_p(s_j) = v_p(m^{j-1}) + v_p(B_j) - v_p(j).$$

We want to show that $v_p(s_j) \rightarrow \infty$ as $j \rightarrow \infty$. This can be achieved by finding the lower bound of $v_p(s_j)$, using a result from [14, Theorem 5.10].

LEMMA 3.5 (von Staudt–Clausen theorem). *Let B_j be a Bernoulli number. Then the fractional part of B_j is given by*

$$\sum_{(p-1)|j} \frac{1}{p}.$$

Suppose $v_p(m) = r$, then we have $v_p(m^{j-1}) = r(j-1)$. By the above lemma, $v_p(B_j) \geq -1$, since B_j contains at most a single factor of p in its denominator. Since $v_p(j) \leq (\log j)/(\log p)$, we obtain

$$v_p(s_j) \geq r(j-1) - \frac{\log j}{\log p} - 1.$$

From this result, it is clear that $v_p(s_j) \rightarrow \infty$ as $j \rightarrow \infty$, and therefore $|s_j|_p \rightarrow 0$, which completes the proof. □

COROLLARY 3.6. *To calculate $\sum_{j \geq 1} (-1)^j (m^{j-1}/a^j)(B_j/j)$ to precision ν , we need to sum up to the smallest j such that*

$$\nu < v_p(m)(j-1) - \frac{\log j}{\log p} - 1. \tag{3.1}$$

To be able to compute using this formula we need to know how many terms of the infinite sum we need to calculate to guarantee correctness up to a given precision.

PROPOSITION 3.7. *For sufficiently large ν , calculating the partial sum of s_j up to $j = (2\nu + 1)/v_p(m) + 1$ provides the correct result modulo p^ν .*

Proof. We need to show that $j = (2\nu + 1)/v_p(m) + 1$ satisfies inequality (3.1). Substituting the value for j , we obtain

$$\begin{aligned} v_p(m)(j-1) - \frac{\log j}{\log p} - 1 - \nu &= \nu - \frac{\log((2\nu + 1)/v_p(m) + 1)}{\log p} \\ &\geq \nu - \frac{\log(2\nu + 2)}{\log p} \quad \text{since } v_p(m) \geq 1 \\ &= \nu - \frac{\log 2 + \log(\nu + 1)}{\log p}. \end{aligned}$$

Consider this as a function in ν . Since it is monotonically increasing for $\nu > 0$, it is positive when $\nu > k$ for some integer k , which satisfies the condition in Corollary 3.6. □

REMARK 3.8. We note that when $p = 2$ and 3 , $k = 3$ and 1 , respectively. For all other primes p , $k \leq 0$, so $j = (2\nu + 1)/v_p(m) + 1$ could be used for almost all values of p .

In practice, one can achieve a better bound on j by solving inequality (3.1) for the particular m , p and ν values.

Possible optimisations to speed up the algorithm include caching common terms in the computation, and performing some computations in \mathbb{Z}_p (most terms in the sum are elements of \mathbb{Z}_p instead of the extension field).

4. Computing R_p

Let K be a number field and U_K its group of units. A system of fundamental units of K is a set of units that form a basis of U_K , modulo torsion. Let u_1, \dots, u_{r+s-1} be such a system. If we fix an embedding from \mathbb{C}_p to \mathbb{C} , then any embedding from K into \mathbb{C}_p can be considered either real or complex, depending on the composite embedding from K to \mathbb{C} . Dirichlet's unit theorem tells us that there are r real embeddings $(\sigma_1, \dots, \sigma_r)$ and s conjugate pairs of complex embeddings $(\sigma_{r+1}, \bar{\sigma}_{r+1}, \dots, \bar{\sigma}_{r+s}, \sigma_{r+s})$. Let δ_i be 1 or 2 when σ_i is respectively real or complex. The p -adic regulator is given by

$$R_p = \det[\delta_j \log_p |\sigma_j(u_i)|]_{1 \leq i, j \leq r+s-1}.$$

R_p is independent of the choice of ordering of the units and embeddings. If K is totally real or CM, then R_p is independent of the choice of embedding from \mathbb{C}_p to \mathbb{C} , but in all other cases of K there may be ambiguities in its definition. Furthermore, it is not clear that we can classify an embedding from \mathbb{C}_p into \mathbb{C} as either real or complex in an efficient manner, and this provided additional reasons to restrict our algorithms to totally real fields.

Thus, for any system of independent units we can easily compute the p -adic regulator from there. We also need the different p -adic embeddings, but these are either trivial to compute using standard techniques for p -adic factorisation or root finding, or otherwise by making use of the \mathbb{Q} -automorphisms of the field and one fixed p -adic embedding. We note that typically the units are not represented with respect to a fixed basis of the field, but as power products $u_i = \prod_{j=1}^r \alpha_j^{e_{i,j}}$ for some (small) elements α_i and some (large) exponents $e_{i,j} \in \mathbb{Z}$. We note that, despite logarithms of power products being trivial to compute, the α_i are not units (although their power product is a unit), and this requires the computation of logarithms of non-units.

To obtain the correct valuation of the p -adic regulator we need a basis for any p -maximal subgroup of the unit group, that is, a subgroup V of the S -unit group U where $p \nmid (V : U)$. In order to obtain such a subgroup, we use the saturation techniques developed in [3], which computes such a group from any subgroup V of full rank. In particular, for abelian fields of moderate conductor, we can obtain such a group from the cyclotomic units of the surrounding cyclotomic field, which allows us to deal with fields of degree too large for the direct computation using class groups.

5. Character selection

Let χ be a Dirichlet character modulo k , which is a multiplicative homomorphism $\chi : (\mathbb{Z}/k\mathbb{Z})^\times \rightarrow \mathbb{C}^\times$.

For any $k|m$, χ also induces a character modulo m . For any given χ the smallest modulus is the conductor of χ , denoted f_χ . Let $\bar{\chi}$ be the conjugate character with the usual definition.

DEFINITION 5.1. Let X be a finite group of Dirichlet characters. Denote by f the lowest common multiple of the conductors of all the characters in X . Let H be the intersection of the kernels of all characters in X , and K the fixed field by H in $\mathbb{Q}[\zeta_f]$. Then X is the set of Dirichlet characters associated with the the field K .

COROLLARY 5.2. X is a subgroup of the characters of $\text{Gal}(\mathbb{Q}[\zeta_f]/\mathbb{Q})$. In fact, X is isomorphic to $\text{Gal}(K/\mathbb{Q})$, and the degree of K/\mathbb{Q} is the order of X .

For each required component in the formula we have already highlighted their computations in the earlier sections. However, we still need to find X to be able to evaluate $L_p(1, \chi)$.

We start by computing the minimal f so that $K \subseteq \mathbb{Q}[\zeta_f]$. If K is already a cyclotomic field, we simply take all even characters of conductor f that are non-trivial. Otherwise, we start with characters of conductor f with order $\deg(K/\mathbb{Q})$. Any further restriction depends on the field in question, in particular the value of f .

If the field is cyclic and f is prime then the characters required are only the primitive ones. However, if f is not prime, then there will be several fields with the same degree and conductor, and thus we must select the characters corresponding to each subfield. Since f , the conductor of K , can be quite large compared to the degree of K , we do not want to compute the f th cyclotomic field explicitly, nor any embeddings from K into $\mathbb{Q}[\zeta_f]$.

ALGORITHM 5.3. *Selecting characters associated to field K*

Input: K

Output: *set of characters X*

- 1: *Set degree and conductor of K to be n and f , respectively*
- 2: *Denote the set of even Dirichlet characters with conductor f and order n by S*
- 3: *Construct map between the ray class group modulo f , Cl_f , and $(\mathbb{Z}/f\mathbb{Z})^\times$*
- 4: *Construct H , the norm group of K , from Cl_f*
- 5: *Let the set of characters in S that act trivially on H be X .*

Proof. We start with $\text{Gal}(\mathbb{Q}[\zeta_f]/\mathbb{Q})$, which is isomorphic to $(\mathbb{Z}/f\mathbb{Z})^\times$. Consider the projection ϕ :

$$\text{Gal}(\mathbb{Q}[\zeta_f]/\mathbb{Q}) \rightarrow \text{Gal}(K/\mathbb{Q}).$$

The kernel of ϕ is $\text{Gal}(\mathbb{Q}[\zeta_f]/K)$, or the automorphisms of $\mathbb{Q}[\zeta_f]$ that fix K . Any character associated to K would act trivially on the kernel.

We know that, from class field theory,

$$\text{Gal}(K/\mathbb{Q}) \sim Cl_f/H,$$

where Cl_f is the ray class field of modulo f [6, Algorithm 4.3.1], and H is the norm group in \mathbb{Z} generated by norms of ideals in K . With knowledge of Cl_f and $\text{Gal}(K/\mathbb{Q})$, we can compute H by taking norms of primes until Cl_f/H reaches the appropriate size. The kernel of ϕ is H , so we need to find the characters that act trivially on H . Although

$$Cl_f \sim (\mathbb{Z}/f\mathbb{Z})^\times$$

and χ acts on $(\mathbb{Z}/f\mathbb{Z})^\times$, there is no canonical map between elements of these two sets. We can construct such a map by examining how various primes map to both Cl_f and $(\mathbb{Z}/f\mathbb{Z})^\times$. This allows us to find the generators of the kernel of ϕ . Since we know the characters required act trivially on the generators, we can test to find the correct characters. □

Since we can compute every part of equation (2.1) except h , we can easily compute h using this formula and find its valuation, which will give the p -part of h .

6. Analysis

We now estimate the complexity of computing $L_p(1, \chi)$, using each of the two methods. We will use classical algorithms for multiplication and division in our comparison.

PROPOSITION 6.1. *The complexity of computation of $L_p(1, \chi)$ using Theorem 2.2 to precision ν is $O(f_\chi \nu^3 d^2)$, where $d = [\mathbb{Q}_p[\zeta_n, \zeta_{f_\chi}] : \mathbb{Q}_p]$.*

Proof. Let d_f be the degree of $\mathbb{Q}_p[\zeta_{f_\chi}]$. Performing each logarithm using classical algorithms to precision ν requires ν calculations, each of complexity $O(d_f^2 \nu^2)$. The remaining multiplication has complexity $O(d^2 \nu^2 \log^2 p)$, giving $O(f_\chi \nu^3 d^2)$ as required. \square

REMARK 6.2. We can estimate what the upper bound on complexity is using only the terms p , f_χ and ν . For fixed f , it follows from Dirichlet's theorem on arithmetic progressions that the primes p for which $p \nmid f$ are equally distributed in $(\mathbb{Z}/f\mathbb{Z})^\times$, and the degree of unramified extension should have the same distribution as orders of elements of $(\mathbb{Z}/f\mathbb{Z})^\times$. This means that the extension degree d is not dependent on the size of p .

For large enough p neither n nor f_χ would contain p as a factor, so it is safe to say we only need an unramified extension, and that $d \leq \phi(f_\chi) < f_\chi$. This gives us an upper bound of $O(f_\chi^3 \nu^3 \log^2 p)$.

PROPOSITION 6.3. *The complexity of computation of $L_p(1, \chi)$ using Theorem 2.3 to precision ν is $O(\text{lcm}(f_\chi, p) d_f^2 \nu^3)$, where $d_f = [\mathbb{Q}_p[\zeta_{f_\chi}] : \mathbb{Q}_p]$.*

Proof. To compute the infinite sum with precision ν requires performing at most $2\nu + 2$ additions, each with complexity in the order of ν^2 , providing a complexity of $O(\nu^3)$ for this sum. The logarithm now has complexity $O(d_f^2 \nu^3)$, and it must be computed for each of the $\text{lcm}(f_\chi, q_p)$ additions in the formula (each of order ν) for a final complexity of $O(\text{lcm}(f_\chi, p) d_f^2 \nu^3)$. \square

REMARK 6.4. The degree of extension depends only on the field degree (the only root of unity required comes from the Dirichlet character, which has order dividing n). Unless $p \mid n$, only an unramified extension is required, and the degree is bounded by $n - 1$. We also have an upper bound on $\text{lcm}(f_\chi, q_p)$ of $2f_\chi p$, so the worst-case scenario for complexity is $O(f_\chi p n^2 \nu^3)$.

It is also unclear what precision is required, although there are numerous heuristic arguments [13] that h is in general small. However, even if we use the Minkowski bound as an upper limit for h , this would mean we need to calculate precision of up to $O(\log(\sqrt{D})/\log p)$.

Comparing the complexities of the two approaches, we see that in addition to the common ν^3 term, the method based on Theorem 2.2 is dependent on $f_\chi d^2$ whilst the approach based on Theorem 2.3 is related to $\text{lcm}(f_\chi, q_p)$. Thus in the case where the degree of the p -adic field constructed is small, the first approach will be faster; the second method would be superior if p is a factor of f_χ . This answers Cohen's question in a remark from [7, p. 304] regarding which is better for computation.

In the classical algorithm to compute the entire class group, the unconditional verification of the computation requires $O(\sqrt{|D|})$ steps [5, Algorithm 6.5.6]. Assuming that we have a tentative class number and unit group for verification, the proposed algorithm requires the computation of approximately f steps. For a number field K with prime degree n , the conductor-discriminant formula reduces to $D = f^{n-1}$. This means that theoretically, the proposed algorithm is asymptotically faster than the existing algorithms for number fields of degree 5 or higher, with improvement increasing for larger n . When only the p -part of the class group is required this would yield a faster computation.

When the degree of the number field is not prime, the relationship between the conductor and discriminant is more complicated. The conductor-discriminant formula [12, Chapter VII, Section 5.11] only states that if X is the set of Dirichlet characters associated to the number field, then

$$D = \prod_{\chi \in X} f_\chi.$$

Since f is the lowest common multiple of all the f_χ , we have $D \leq f^{n-1}$. We would need an upper bound to show that the proposed algorithm is an improvement.

For number fields with Galois group $C_2 \times C_2$ we have either $D = f^2$ or $D = 4f^2$, and we do not expect any improvement over existing algorithms. In number fields with Galois group C_4 the relationship between D and f is more complicated, with either $D = f^3$ or $D = gf^2$, where g is a divisor of f . Moving up to fields with Galois group C_6 , it becomes $D = g_1g_2^2f$, where both g_1 and g_2 are divisors of f . The lack of relationship between g_1 and g_2 makes it difficult to compare the algorithms.

Instead of investigating the individual relationship between D and f for number fields with a particular Galois group, we provide an asymptotic limit on what the ratio would be, based on what we know about the densities of discriminants and conductors.

Let $\mathcal{D}(X)$ and $\mathcal{F}(X)$ be the number of algebraic number fields (with fixed abelian Galois group G) contained in some fixed algebraic closure, with discriminant up to X and conductor up to X , respectively. Based on [11], for any $\varepsilon > 0$ we have the formula

$$\mathcal{F}(X) = XP(\log X) + O(X^{1-3/(v_0+6)+\varepsilon}),$$

where $P(\log X)$ is a polynomial (in $\log X$) of degree d_0 and leading coefficient c , d_0 and v_0 being constants dependent on G . This states that the number of fields with conductor at most X is asymptotic to $cX \log^{d_0} X$, for some constant c .

Wright [15] showed that as $X \rightarrow \infty$,

$$\mathcal{D}(X) \approx \frac{c(G)\alpha}{(v-1)!} X^{1/\alpha} \log^{v-1} X.$$

Here $c(G)$ is a constant depending on G , $\alpha = |G|(1 - 1/Q)$, where Q is the smallest prime divisor of the order of G , and $v = |G|_Q/Q$, where $|G|_Q$ is the number of elements with order Q .


While we know that the discriminant does not grow monotonically with the conductor, nevertheless we do have $D \rightarrow \infty$ as $f \rightarrow \infty$. Therefore, if we have $\mathcal{D}(X) = \mathcal{F}(Y)$, then we know that a field with conductor Y will have discriminant X (at least asymptotically). If we can solve for X explicitly then we are able obtain a relationship between f and D (and use this to calculate the ratio).

The dominant term in each of $\mathcal{D}(X)$ and $\mathcal{F}(Y)$ is the power of X or Y , respectively. Since we are only interested in asymptotic behaviour, we will only consider these. Since $X^{1/\alpha} \approx Y$, we find that

$$\frac{\log |D|}{\log f} \rightarrow \alpha (= |G|(1 - 1/Q)).$$

Since α is dependent only on the group order, all fields of the same degree would have the same asymptotic ratio, regardless of the actual structure of the Galois group. For fields with composite degree greater than 4, we have $\alpha \geq 3$, so we can conclusively say that the proposed algorithms offer an advantage on all fields of degree 5 or higher over the classical method.

REMARK 6.5. We have not taken into consideration the complexity of computing R_p , which is required for computing p -part of h , due to the fact that we are unable to provide a full complexity analysis. Furthermore, the same algorithm for saturation would be used to speed up the existing algorithm, albeit at multiple primes rather than a single prime in our case. Partial analysis of the saturation algorithm shows that no part of the algorithm is worse than polynomial in either p or n , which would indicate perhaps $O(p^k n^l)$, which is not inconsistent with results provided by heuristic arguments that the overall algorithm remains linear in f . For example, for fields with Galois group C_7 , linear complexity is clearly observed in the graph of conductor and complexity in Figure 1.

FIGURE 1. Computation time for fields with Galois group C_7 .

7. Examples

We highlight three specific examples here. The first is a complete worked example, while the latter are instances where p -adic verification can be achieved in less time than the classical approach (using either a GRH or unconditional bound). These examples were performed using Magma v2.19 [4] (note that the default bound PARI uses assumes the GRH).

7.1. $\mathbb{Q}[\sqrt{40}]$, $p = 2$

Neither method from [1] or [9] can deal with this example ($p = 2$ and p divides the field degree).

Both D and f are 40. Using the Iwasawa approach, $40 = 2^3 \cdot 5$ and $2^4 = 1 \pmod{5}$, so we construct the unramified extension of \mathbb{Q}_2 defined by the polynomial $x^4 + x + 1$. Let α be a root of the defining polynomial in this extension. We then need to construct two ramified extensions of this field, first with the polynomial $x + 2$, followed by $x^4 + 4x^3 + 6x^2 + 4x + 2$. Let β be the root of $x^4 + 4x^3 + 6x^2 + 4x + 2$ in the final extension field.

We obtain an approximation to the 40th root of unity (correct up to modulo 2^5), which is required for later calculations, as

$$(-13\alpha^3 + 8\alpha^2 + 14\alpha - 2)\beta - 13\alpha^3 + 8\alpha^2 + 14\alpha - 2.$$

The characters required are of order 2, with conductor 40. It turns out that only a single character χ is required, with $\chi(17) = -1$, $\chi(21) = -1$ and $\chi(31) = 1$.

The p -adic zeta function and p -adic regulator are found to be β^4 and γ (a root to the equation $x^2 - 10$ in \mathbb{Q}_2), with valuation 4 and 1, respectively. Putting this into the p -adic class number formula, we get $v_2(h) = 1$ as required.

7.2. $\mathbb{Q}[\theta]$, $p = 2$

Take the field of \mathbb{Q} adjoined by the root of $x^7 - x^6 - 354x^5 - 979x^4 + 30030x^3 + 111552x^2 - 715705x - 2921075$. The conductor is 827 and the Minkowski bound is 3461471. The classical method provides a conditional class group of $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ in less than 1 second, but requires another 160 s to verify the result. 2-adic verification based on computing hR_p takes 1.5 s, significantly faster than the classical approach.

7.3. $\mathbb{Q}[\phi]$, $p = 11$

Here, ϕ is a root of the polynomial

$$x^{11} - x^{10} - 930x^9 - 1049x^8 + 254577x^7 - 177105x^6 - 28898705x^5 \\ + 105363794x^4 + 1065225462x^3 - 7828574944x^2 + 15893036840x - 7589985325.$$

The conductor is 2047 and the Minkowski bound is 5028348788074. The classical method is capable of computing a tentative class group of C_{11} for this field in approximately 13 s, and another 61 s to check up to the Bach bound (69752). It would be infeasible to verify this result using the Minkowski bound.

11-adic verification of the class group takes 4.5 s. Furthermore, a complete p -adic calculation determines the p -valuation of the class number to be 1, with the entire process taking 22 s. This is even faster than using the Bach bound and represents an improvement over the existing algorithm, even if GRH is assumed.

References

1. M. AOKI and T. FUKUDA, 'An algorithm for computing p -class groups of abelian number fields', *Algorithmic number theory*, Lecture Notes in Computer Science 4076 (Springer, Berlin, 2006) 56–71.
2. D. BERNSTEIN, 'Fast multiplication and its applications', *Algorithmic number theory: lattices, number fields, curves and cryptography*, Mathematical Sciences Research Institute Publications 44 (Cambridge University Press, Cambridge, 2008) 325–384.
3. J.-F. BIASSE and C. FIEKER, 'New techniques for computing the ideal class group and a system of fundamental units in number theory', *ANTS X: Proceedings of the Tenth Algorithmic Number Theory Symposium* (Mathematical Sciences Publishers, Berkeley, CA, 2013).
4. W. BOSMA, J. CANNON and C. PLAYOUST, 'The Magma algebra system. I. The user language', *J. Symbolic Comput.* 24 (1997) 235–265.
5. H. COHEN, *A course in computational algebraic number theory*, Graduate Texts in Mathematics 138 (Springer, Berlin, 1993).
6. H. COHEN, *Advanced topics in computational number theory*, Graduate Texts in Mathematics 193 (Springer, New York, 2000).
7. H. COHEN, *Number theory. Vol. II. Analytic and modern tools*, Graduate Texts in Mathematics 240 (Springer, New York, 2007).
8. G. GRAS and M.-N. GRAS, 'Calcul du nombre de classes et des unités des extensions abéliennes réelles de \mathbb{Q} ', *Bull. Sci. Math.* (2) 101 (1977) no. 2, 97–129.
9. T. HAKKARAINEN, 'On the computation of class numbers of real abelian fields', *Math. Comp.* 78 (2009) no. 265, 555–573.
10. K. IWASAWA, *Lectures on p -adic L -functions*, Annals of Mathematics Studies 74 (Princeton University Press/University of Tokyo Press, Princeton, NJ/Tokyo, 1972).
11. S. MÁKI, 'The conductor density of abelian number fields', *J. Lond. Math. Soc.* (2) 47 (1993) no. 1, 18–30.
12. J. NEUKIRCH, *Algebraic number theory*, Grundlehren der Mathematischen Wissenschaften 322 (Springer, Berlin, 1999); translated from the 1992 German original and with a note by Norbert Schappacher. With a foreword by G Harder.
13. R. SCHOOF, 'Class numbers of real cyclotomic fields of prime conductor', *Math. Comp.* 72 (2003) no. 242, 913–937.
14. L. C. WASHINGTON, *Introduction to cyclotomic fields*, 2nd edn. Graduate Texts in Mathematics 83 (Springer, New York, 1997).
15. D. J. WRIGHT, 'Distribution of discriminants of abelian extensions', *Proc. Lond. Math. Soc.* (3) 58 (1989) no. 1, 17–50.

Claus Fieker
 Fachbereich Mathematik
 Technische Universität Kaiserslautern
 Germany

fieker@mathematik.uni-kl.de

Yinan Zhang
 School of Mathematics and Statistics
 The University of Sydney
 Australia

y.zhang@sydney.edu.au