



Universiteit
Leiden
The Netherlands

Random walks on Arakelov class groups

Boer, K. de

Citation

Boer, K. de. (2022, September 22). *Random walks on Arakelov class groups*. Retrieved from <https://hdl.handle.net/1887/3463719>

Version: Publisher's Version

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

Downloaded from: <https://hdl.handle.net/1887/3463719>

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

7. The Power Residue Symbol is in ZPP

7.1. Summary

In this chapter we show that, assuming the Riemann hypothesis for Hecke L-functions on the cyclotomic fields $\mathbb{Q}(\zeta_m)$, the problem of computing the m -th power residue symbol in a field containing the m -th root of unity lies in the complexity class ZPP. In other words, there exists an algorithm that computes power residue symbols within probabilistic polynomial time in the input size. Though this algorithm never outputs an incorrect output, it might simply give *no* output with a certain constant probability. The probability here is over, say, random coin flips, which allows the algorithm to repeat until having a negligible error probability. Such algorithms are also known as Las Vegas algorithms.

The proof of the polynomial running time consists of essentially two parts, which are treated separately in Section 7.4 and Section 7.5. The former part consists of an efficient reduction from general power residue symbols to power residue symbols in cyclotomic fields; this reduction is due to Lenstra [Len95] and Squirrel [Squ97]. The latter part is a new result and consists of a proof that power residue symbols in cyclotomic fields can be computed efficiently, assuming the Extended Riemann Hypothesis for Hecke L-functions on cyclotomic fields. The key ingredient for this algorithm to be provable is the sampling algorithm of the previous Chapter 6. By combining these two parts, one obtains a conditional proof that power residue symbols can be computed efficiently in any number field.

7.2. Introduction

The power residue symbol often plays a significant role in algorithms in which *residuosity* is involved, which is about distinguishing m -th powers from non- m -th powers modulo an ideal in a number field. In such case, the m -th power residue symbol serves as a first check, as it should be equal to one in the case of an m -th power.

Examples of cryptographic schemes involving residuosity and that need a fast computation of power residue symbols include [SW95; GM84; Sch98; Wil85], which mostly consider m being prime and below 12. It should be noted that these cryptographic schemes (and actually, most residuosity-based schemes) are not quantum secure, due to their susceptibility to Shor's efficient quantum algorithm for factoring [Sho94]. In fact, if one is allowed to use a quantum computer, a very simple algorithm for the power residue exists, by just factoring the bottom input ideal of $(\frac{\alpha}{b})$. So, to be clear, in this chapter we will solely consider classical computing power.

Efficient algorithms for the m -th power residue symbol for specific small cases of $m \leq 11$ are studied extensively [CS10; DF05; Wei02; Wil85; SW95; Lem00]. A first attempt to design an efficient algorithm for *general* m -th power residue symbols (i.e., for all m) was done by Squirrel in his undergraduate thesis [Squ97]. In that work, Squirrel derives an efficient reduction from power residue symbols in general number fields to those in cyclotomic fields based on an idea of Lenstra [Len95]. Squirrel also proposes an algorithm for computing power residue symbols in cyclotomic fields, but it relies on heavy precomputations and is therefore not polynomial for varying m [Squ97, Ch. 5, §3]. On top of that, the algorithm also seems unfeasible in terms of practical running time.

Later, an algorithm for m -th power residue symbols that seems practically feasible and runs heuristically in polynomial time (for varying m) was given by the author of this PhD thesis [Boe16; BP17]. In this chapter we prove that a variant of this heuristic algorithm lies in the complexity class ZPP, assuming the Extended Riemann Hypothesis for Hecke L-functions on cyclotomic fields.

It should be noted that the aforementioned algorithms tailored to specific small $m \leq 11$ are by far more efficient than this more general algorithm, are mostly deterministic and also do not require any variant of the Riemann hypothesis.

Difference between the power residue symbol algorithm of this chapter and the heuristic algorithm in [Boe16; BP17]

The key difference between the power residue symbol algorithm of this chapter and that of [Boe16; BP17] is their *purpose*. The algorithm described in this chapter is namely specifically constructed in such a way that the proof of its polynomial time complexity is as simple as possible. The heuristic algorithm in [Boe16; BP17], however, is much more directed toward implementation and a fast practical running time (for an implementation, see [Boe17]). This distinction in purpose lead to the following key differences between the two algorithms.

The provable algorithm does not use the Hilbert reciprocity law. As opposed to the heuristic algorithm, the provable algorithm of this chapter does *not* use *Hilbert reciprocity*. In other words, the following reciprocity law involving Hilbert symbols does not play any role in the provable algorithm of this chapter.

$$\left(\frac{\alpha}{\beta}\right)_m \left(\frac{\beta}{\alpha}\right)_m^{-1} = \prod_{\mathfrak{p}|m\infty} (\alpha, \beta)_{\mathfrak{p}},$$

Here, $(\alpha, \beta)_{\mathfrak{p}}$ are the m -th Hilbert symbols in the completion $\mathbb{Q}(\zeta_m)_{\mathfrak{p}}$ (e.g., [Neu85, Ch. III, §5 and Ch. IV, §9]). Avoiding the Hilbert reciprocity law has as an advantage that there is no need to compute Hilbert symbols in the provable algorithm. In the heuristic algorithm, the computation of such Hilbert symbols relied on a efficient and provable algorithm of Bouw [Bou21].

The provable algorithm uses the Artin reciprocity law. Instead, the provable algorithm of this chapter uses a different reciprocity law, namely the *Artin reciprocity* law (see Lemma 7.3), which states that for elements $\kappa \in K^{m,1}$ in a specific ray, the power residue symbol $\left(\frac{\alpha}{\kappa}\right) = 1$ for all $\alpha \in K^*$. This turned out to be easier to use in a proof and has as an additional advantage that no computation of Hilbert symbols is needed. In fact, one can even use this provable algorithm to compute Hilbert symbols instead (see Section 7.6.1).

The provable algorithm does not use LLL-reduction. The heuristic algorithm of [Boe16; BP17] uses LLL-reduction to minimize sizes of the input while this is omitted in the provable algorithm for the sake of brevity and provability.

7.3. Preliminaries

In this chapter, K is a degree $n = [K : \mathbb{Q}]$ number field containing the m -th cyclotomic number field, i.e., $K \supseteq \mathbb{Q}(\zeta_m)$, where ζ_m is a primitive m -th root of unity. The main subject of this chapter is the *power residue symbol*, a map that partially captures m -th residuosity.

This power residue symbol takes as an input an ideal \mathfrak{b} in an order R of K and an element $\alpha \in R$, and outputs an m -th root of unity ζ_m^k . The symbol and its definition resembles that of the Jacobi symbol, for example in the sense that it can be defined in terms of prime ideals first, and can subsequently be multiplicatively extended to general ideals.

Definition 7.1 (Power residue symbol). *Let $\mathfrak{p} \nmid m$ be a prime ideal in an order R of $K \ni \zeta_m$ and let $\alpha \in R$ be an element coprime with m and \mathfrak{p} . We define $\left(\frac{\alpha}{\mathfrak{p}}\right) \in \langle \zeta_m \rangle = \{\zeta_m^k \mid k \in \mathbb{N}\}$ to be the m -th root of unity that satisfies*

$$\left(\frac{\alpha}{\mathfrak{p}}\right) \equiv \alpha^{\frac{N(\mathfrak{p})-1}{m}} \pmod{\mathfrak{p}}.$$

For general ideals \mathfrak{b} in R coprime with m we then use the prime ideal factorization $\mathfrak{b} = \prod_j \mathfrak{p}_j^{e_j}$ to define the power residue symbol $(\frac{\alpha}{\mathfrak{b}})$.

$$\left(\frac{\alpha}{\mathfrak{b}}\right) := \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right)^{e_j}.$$

By the very definition of the power residue symbol ‘above’ prime ideals, they can be computed efficiently and deterministically.

Lemma 7.2. *Let $\mathfrak{p} \subseteq \mathbb{Z}[\zeta_m]$ be a prime ideal not dividing m . Then the power residue symbol $(\frac{\alpha}{\mathfrak{p}})$ can be computed within $\text{poly}(m, \log \mathcal{N}(\mathfrak{p}), \log |\mathcal{N}(\alpha)|)$ time.*

Proof. By the power residue symbol formula for prime ideals we have $(\frac{\alpha}{\mathfrak{p}}) \equiv \alpha^{(\mathcal{N}(\mathfrak{p})-1)/m}$ modulo \mathfrak{p} . We compute the (modular) Hermite normal form [SL96; HM91] [Coh93, §2.4.2] of the ideal \mathfrak{p} , which allows to have a unique representative for each element in $\mathbb{Z}[\zeta_m]/\mathfrak{p}$. By modular exponentiation, can compute $\alpha^{(\mathcal{N}(\mathfrak{p})-1)/m}$ modulo \mathfrak{p} within time $\text{poly}(m, \log \mathcal{N}(\mathfrak{p}), \log |\mathcal{N}(\alpha)|)$. \square

The following lemma shows that the power residue symbol is trivial for certain values of the lower input. Specifically, considering a fixed upper input for the power residue symbol, the map $(\frac{\alpha}{\cdot}) : K \rightarrow \langle \zeta_m \rangle$ has a kernel that includes the ray $K^{\mathfrak{m},1}$ with $\mathfrak{m} = m^m \cdot \alpha$. This particular fact forms one of the very key ingredients of the efficient power residue symbol algorithm.

Lemma 7.3. *For all $\alpha \in \mathbb{Z}[\zeta_m]$ coprime with m , and all $\kappa \in \mathbb{Q}(\zeta_m)^*$ with $\text{ord}_{\mathfrak{p}}(\kappa) \geq 0$ for all $\mathfrak{p} | \alpha m$, we have,*

$$\left(\frac{\alpha}{1 + \kappa \cdot m^m \cdot \alpha}\right) = 1$$

Proof. Denote $K = \mathbb{Q}(\zeta_m)$ and $L = \mathbb{Q}(\zeta_m, \sqrt[m]{\alpha})$ for $\alpha \in \mathbb{Q}(\zeta_m)$. The power residue symbol $\left(\frac{\alpha}{\mathfrak{p}}\right) \in \langle \zeta_m \rangle$ in $\mathbb{Q}(\zeta_m)$ has the following relation with the Artin symbol [Lem00, §4.1] [Koc97, Ch. 2, §2.1]

$$\left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \sqrt[m]{\alpha} = \left(\frac{\mathfrak{b}}{L/K}\right) [\sqrt[m]{\alpha}].$$

Denote $\mathfrak{f}_{L/K}$ for the *conductor* of the extension L/K . For any modulus \mathfrak{m} satisfying $\mathfrak{f}_{L/K} \mid \mathfrak{m}$, the *kernel* of the Artin symbol $\left(\frac{\cdot}{L/K}\right) : \mathcal{I}_K^{\mathfrak{m}} \rightarrow G$ contains the *ray* $K^{\mathfrak{m},1}$, the multiplicative subgroup of K^* generated by elements $\kappa \in \mathbb{Z}[\zeta_m]$ that are 1 modulo \mathfrak{m} . This is a consequence of the Artin reciprocity law [Chi08, Thm. 2.1].

It remains to show that $\mathfrak{m} = m^m \alpha$ satisfies $\mathfrak{f}_{L/K} \mid \mathfrak{m}$, i.e., that $\mathfrak{f}_{L/K} \mid m^m \alpha$. If we can prove that fact, the result follows, since $1 + k \cdot m^m \cdot \alpha \in K^{\mathfrak{m},1}$ for κ satisfying $\text{ord}_{\mathfrak{p}}(\kappa) \geq 0$ for all $\mathfrak{p} \mid \alpha m$.

In the following, we prove that $\mathfrak{f}_{L/K} \mid (m^m \alpha)$. Since α is required to be coprime with m , and the degree of the extension satisfies $[L : K] \mid m$, any $\mathfrak{p} \mid (\alpha)$ is tamely ramified in the extension L/K , because $\mathfrak{p} \nmid m$. Therefore, we have, [CG05, Ch. 2, Prop. 1.6.3] [CS08, Eq. (3.10) and Eq. (3.11)]

$$\text{For all } \mathfrak{p} \mid (\alpha) : \text{ord}_{\mathfrak{p}}(\mathfrak{f}_{L/K}) = 1.$$

From the same results, or from the fact that $\mathfrak{f}_{L/K} \mid \Delta_{L/K} \mid m^m \alpha^{m-1}$ [CF10, Lm. 5, Ch. 3] [NS13, Ch. VII, Prop. 11.9] follows that $\mathfrak{f}_{L/K} \mid (m^m \alpha)$. \square

The following result, namely, multiplicativity in the bottom input of the power residue symbol, can be found in [Neu85, Ch. 4, Eq. (9.2)] or [Koc97, Thm. 2.13].

Lemma 7.4. *Let K be a number field containing $\mathbb{Q}(\zeta_m)$. Let $\mathfrak{b}, \mathfrak{c} \in \mathcal{I}_K$ be coprime with m . For all $\alpha \in K$ coprime with $\mathfrak{b}, \mathfrak{c}$ and m , we have*

$$\left(\frac{\alpha}{\mathfrak{b}\mathfrak{c}}\right) = \left(\frac{\alpha}{\mathfrak{b}}\right) \left(\frac{\alpha}{\mathfrak{c}}\right).$$

The last important lemma of this preliminaries concerns the local density of the prime ideals coprime to \mathfrak{m} . It turns out that for large enough r and not too large modulus \mathfrak{m} , the density $\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n]$ does not differ so much from the density $\delta_{\mathcal{P}}[r^n]$ of all prime ideals in a number field. This density is known to be close to $\frac{1}{\rho_K \cdot \log(r^n)}$, where ρ_K is the residue of the Dedekind zeta function $\zeta_K(s)$ at the pole at $s = 1$.

This density is important because it is tightly related to the success probability of the power residue symbol algorithm of this chapter. This is because the power residue symbol algorithm involves prime ideal sampling, as in Chapter 6.

Lemma 7.5. *Let $\mathcal{P}^{\mathfrak{m}} = \{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \mid \mathfrak{p} \text{ prime}\}$ and let $\omega(\mathfrak{m})$ denote the number of different prime ideal divisors of \mathfrak{m} . Then, for all $r^n \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11}, 16 \cdot \omega(\mathfrak{m})^2)$, we have*

$$\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n] \geq \frac{1}{4n \cdot \rho_K \cdot \log r}.$$

Recall that $n = [K : \mathbb{Q}]$, the degree of the number field K .

Proof. By Lemma 2.13, considering $x \in [(r/e)^n, r^n]$ and Definition 6.6, we have

$$\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n] = \min_{x \in [(r/e)^n, r^n]} \frac{\pi_K^{\mathfrak{m}}(x)}{\rho_K x} \geq \frac{x / \log x}{4\rho_K x} \geq \frac{1}{4n\rho_K \log(r/e)} \geq \frac{1}{4n\rho_K \log r}.$$

□

7.4. Reduction to Cyclotomic Fields

7.4.1. Introduction

In this section, we will show that the computation of the m -th power residue symbol in any order R (of a number field) containing $\mathbb{Z}[\zeta_m]$ reduces to the

computation of (polynomially) many power residue symbols in $\mathbb{Z}[\zeta_m]$, the ring of integers of the m -th cyclotomic field.

The strategy of this proof is described in a paper of Lenstra [Len95], in which the special case $m = 2$ is elaborately worked out. For general $m > 2$, a full description of this reduction is given by Squirrel in his undergraduate thesis [Squ97]. In this section we will follow closely the reasoning of Squirrel and Lenstra, omitting precise complexity claims; any of the steps in this reduction runs in time polynomial in the input size.

In the following section, we give an overview of the proof of this reduction, postponing the definitions and proofs to a later moment.

7.4.2. Proof Strategy

Introduction. In this proof summary, we will consider number fields K containing all m -th roots of unity, i.e. $K \supseteq \mathbb{Q}(\zeta_m)$. Instead of the maximal order \mathcal{O}_K , which might be very hard to compute, we will mainly consider general orders $R \subseteq \mathcal{O}_K$ of K .

The main purpose of this proof overview is to show on a high level that we can reduce the computation of the power residue symbol $(\frac{\alpha}{\mathfrak{b}})_{m,K}$ for an element $\alpha \in R$ and an ideal $\mathfrak{b} \subseteq R$ to the computation of power residue symbols in the cyclotomic field $\mathbb{Q}(\zeta_m)$.

Signature identity. The power residue symbol $(\frac{\alpha}{\mathfrak{b}})_{m,K}$ is equal to another special quantity, $(m_\alpha, R/\mathfrak{b})$, which we will call the *signature*. This signature captures certain behavior of the multiplication map $m_\alpha : x \mapsto \alpha \cdot x$ on the finite $\mathbb{Z}[\zeta_m]$ -module R/\mathfrak{b} . Because of this equality, we can shift our attention to computing the signature $(m_\alpha, R/\mathfrak{b})$.

Invariant factor decomposition of R/\mathfrak{b} . A very important observation is the fact that the signature $(m_\alpha, R/\mathfrak{b})$ only depends on the structure of R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module. Using an analogue of the invariant factor decomposition

for finite modules over Dedekind domains (see [Coh99, Thm. 1.2.30]), we can efficiently compute a decomposition of R/\mathfrak{b} of the following shape.

$$R/\mathfrak{b} = \gamma_1 \mathbb{Z}[\zeta_m]/\mathfrak{d}_1 \oplus \cdots \oplus \gamma_k \mathbb{Z}[\zeta_m]/\mathfrak{d}_k, \quad (7.107)$$

where $\gamma \in R$ and \mathfrak{d}_j are ideals of $\mathbb{Z}[\zeta_m]$ that satisfy $\mathfrak{d}_j = \prod_{i=1}^j \mathfrak{c}_i$ for ideals \mathfrak{c}_i of $\mathbb{Z}[\zeta_m]$ that are neither the zero or the unit ideal. In other words, $\mathfrak{d}_{j+1}/\mathfrak{d}_j = \mathfrak{c}_{j+1}$ for $j \in \{1, \dots, k-1\}$ and $\mathfrak{d}_1 = \mathfrak{c}_1$. This computation shows that we can shift our focus to modules of a form as described in Equation (7.107).

The signature is compatible with short exact sequences. Let M', M, M'' be $\mathbb{Z}[\zeta_m]$ -modules with respective ($\mathbb{Z}[\zeta_m]$ -module compatible) automorphisms ϕ', ϕ and ϕ'' , that fit into the following commuting diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' & \longrightarrow & 0 \\ & & \downarrow \phi' & & \downarrow \phi & & \downarrow \phi'' & & \\ 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' & \longrightarrow & 0 \end{array}$$

Then we have $(\phi, M) = (\phi', M') \cdot (\phi'', M'')$, i.e., the signature of the ‘middle’ module can be computed with the signatures of the ‘outer’ modules.

The determinant formula. For $\mathbb{Z}[\zeta_m]$ -modules isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ for some $t \in \mathbb{Z}_{>0}$ and integral ideal \mathfrak{c} of $\mathbb{Z}[\zeta_m]$, we can compute the signature by means of the determinant formula. Any automorphism ϕ of $(\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ can be described by a non-degenerate matrix with entries in $\mathbb{Z}[\zeta_m]/\mathfrak{c}$, which makes $\det(\phi) \in \mathbb{Z}[\zeta_m]/\mathfrak{c}$ a well-defined quantity. The determinant formula then reads as follows.

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = \left(\frac{\det(\phi)}{\mathfrak{c}} \right)_{m, \mathbb{Q}(\zeta_m)}. \quad (7.108)$$

Note that this reduces the computation of this specific signature to a power residue symbol in the cyclotomic field $\mathbb{Q}(\zeta_m)$.

Applying induction on the components of the module. Denoting $M = R/\mathfrak{b}$, we have the following exact sequence

$$0 \rightarrow M/(\mathfrak{c}_1 M) \rightarrow M \rightarrow \mathfrak{c}_1 M \rightarrow 0,$$

where \mathfrak{c}_1 is the first factor in the invariant factor decomposition (Equation (7.107)). Because of the compatibility of the signature with short exact sequences, it is enough to compute the signatures $(\phi_\alpha, M/(\mathfrak{c}_1 M))$ and $(\phi_\alpha, \mathfrak{c}_1 M)$.

The first module, $M/(\mathfrak{c}_1 M)$, can be shown to be isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c}_1)^k$, and therefore the determinant formula applies (see Equation (7.108)).

The last module, $\mathfrak{c}_1 M$, can be shown to have less ‘components’ than M itself; $k - 1$ instead of k .

$$\mathfrak{c}_1 M = \bigoplus_{j=1}^{k-1} \gamma_j \mathbb{Z}[\zeta_m]/\tilde{\mathfrak{d}}_j,$$

where $\tilde{\mathfrak{d}}_j = \mathfrak{d}_j/\mathfrak{c}_1$, and where \mathfrak{d}_j are obtained from the invariant factor decomposition of $M = R/\mathfrak{b}$.

Conclusion. By induction, we can therefore conclude that the computation of $(\phi_\alpha, R/\mathfrak{b})$ reduces to k power residue symbols $\left(\frac{d_j}{c_j}\right)_{m, \mathbb{Q}(\zeta_m)}$ for $j \in \{1, \dots, k\}$ in the cyclotomic field $\mathbb{Q}(\zeta_m)$. Here, \mathfrak{c}_j are the invariant factors of the module R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module and $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are determinants of associated automorphisms.

7.4.3. Signature Identity

Definition 7.6 (Admissible modules). *We call a $\mathbb{Z}[\zeta_m]$ -module M admissible if $|M|$ is finite and $\gcd(|M|, m) = 1$.*

Letting the group $\langle \zeta_m \rangle = \{\zeta_m^j \mid j \in \mathbb{Z}/m\mathbb{Z}\}$ act on an admissible module M , we can directly deduce that this action must be free on $M \setminus 0$. Namely, suppose

that there exists an $x \in M$ with $\zeta_m^j x = x$. Then we have $(\zeta_m^j - 1)x = 0$, which implies $mx = 0$ (as $(\zeta_m^j - 1) \mid m$). Since $|M|x = 0$, $mx = 0$ and $\gcd(|M|, m) = 1$, we have $1 \cdot x = 0$.

This means that $M \setminus 0$ can be written as a disjoint union of orbits $\langle \zeta_m \rangle \cdot x$ (for some $x \in M$), where the orbits have precisely m elements. This directly implies $|M| = tm + 1$, where t is the number of orbits in $M \setminus 0$. Summarizing, any admissible module M satisfies $|M| \equiv 1$ modulo m .

Let M be an admissible $\mathbb{Z}[\zeta_m]$ -module and let $\phi : M \rightarrow M$ be a bijective function satisfying $\phi(\zeta_m \cdot x) = \zeta_m \cdot \phi(x)$ for all $x \in M$. Then ϕ acts faithfully on the $\langle \zeta_m \rangle$ -orbits of M , as $\phi(\langle \zeta_m \rangle \cdot x) = \langle \zeta_m \rangle \cdot \phi(x)$. In other words, ϕ induces a permutation on the quotient set $M/\langle \zeta_m \rangle$, fixing $0 \in M$.

Example 7.7. Put $K = \mathbb{Q}(\zeta_6, \sqrt[3]{2})$, a degree 6 extension of \mathbb{Q} . The subring $R = \mathbb{Z}[\zeta_6, \sqrt[3]{2}]$ is an order in K which has the following R -ideal $\mathfrak{p}_5 = (5, 3 - \sqrt[3]{2})$. Then the $\mathbb{Z}[\zeta_6]$ -module R/\mathfrak{p}_5 has 25 elements; one of them is zero, and the others fall into four $\langle \zeta_6 \rangle$ -orbits of length six, see Figure 7.1.




Figure 7.1.: The multiplicative action of $\langle \zeta \rangle$ on the 25 elements of R/\mathfrak{p}_5 as in Example 7.7, where $\zeta = \zeta_6$, a 6-th primitive root of unity. It consists of one zero-orbit of length one, and four orbits of length 6.

Let $S \subseteq M$ be a representative set for $M/\langle \zeta_m \rangle$, i.e., $M = \bigcup_{s \in S} \langle \zeta_m \rangle s$ (where the union is disjoint). Then, the action of ϕ on $M/\langle \zeta_m \rangle$ induces a bijection $s \mapsto s^\phi$ on S . Here $s^\phi \in S$ is the unique representative in S satisfying $\phi(\langle \zeta_m \rangle s) = \langle \zeta_m \rangle s^\phi$. Note that this means that $\phi(s) \in \langle \zeta_m \rangle s^\phi$, making the fraction $\frac{\phi(s)}{s^\phi} \in \langle \zeta_m \rangle$ well-defined for all $s \in S \setminus 0$. We then arrive at the following definition.

Definition 7.8 (Signature). Let M be an admissible $\mathbb{Z}[\zeta_m]$ -module, let $\phi : M \rightarrow M$ be $\mathbb{Z}[\zeta_m]$ -module homomorphism and let $S \subseteq M$ be a representative

7. The Power Residue Symbol is in ZPP

set for $M/\langle \zeta_m \rangle$. Then we define the signature $(\phi, M) \in \langle \zeta_m \rangle$ as follows.

$$(\phi, M) = \prod_{s \in S \setminus 0} \frac{\phi(s)}{s^\phi} \quad (7.109)$$




Figure 7.2.: The signature of a map ϕ forgets about the permutation of the $\langle \zeta_m \rangle$ -orbits. Instead, it captures the ‘compound deviation’ of the images of representatives from the representative of the orbits that image lives in. For example, the ϕ -image of the green dot deviates $+1$ from the yellow representative in its orbit.

Remark 7.9. The definition above can be generalized to any bijective map $M \rightarrow M$ that commutes with ζ_m [Squ97], but for our purposes it is enough to consider $\mathbb{Z}[\zeta_m]$ -module homomorphisms.

The very nature of the definition shows that (ϕ, M) does not depend on the choice of the representative set S . Namely, changing a single $s \in S$ into $s' = \zeta_m^j \cdot s$ causes a ζ_m^j to appear once in the numerator of a factor in Equation (7.109) and once in the denominator of a factor in Equation (7.109); therefore it does not change the overall value.

Lemma 7.10. Let R be an order in a number field K with $\mathbb{Z}[\zeta_m] \subseteq R$. Let \mathfrak{p} be a prime ideal in R , coprime with m . Let $\alpha \in R$ such that $\bar{\alpha} = \alpha \bmod \mathfrak{p} \in (R/\mathfrak{p})^*$ and denote $\phi_\alpha : R/\mathfrak{p} \rightarrow R/\mathfrak{p}, x \mapsto \bar{\alpha} \cdot x$. Then

$$\left(\frac{\alpha}{\mathfrak{p}} \right)_m = (\phi_\alpha, R/\mathfrak{p})$$

Proof. Taking a representative set S for $M = R/\mathfrak{p}$ (modulo $\langle \zeta_m \bmod \mathfrak{p} \rangle$) we write out the definition of $(\phi_\alpha, R/\mathfrak{p})$ (see Definition 7.8). In the following chain of equalities we make use of the fact that $M = R/\mathfrak{p}$ (next to a $\mathbb{Z}[\zeta_m]$ -module) is also a field, so that division and multiplication of elements there make sense.

$$(\phi_\alpha, R/\mathfrak{p}) = \prod_{s \in S \setminus 0} \frac{\phi_\alpha(s)}{s^{\phi_\alpha}} = \prod_{s \in S \setminus 0} \frac{\bar{\alpha} \cdot s}{s^{\phi_\alpha}} = \bar{\alpha}^{|S \setminus 0|} \frac{\prod_{s \in S \setminus 0} s}{\prod_{s \in S \setminus 0} s^{\phi_\alpha}} = \bar{\alpha}^{|S \setminus 0|}.$$

The last inequality follows from the fact that $s \mapsto s^{\phi_\alpha}$ is a bijection on $S \setminus 0$. As $|S \setminus 0| = \frac{|M|-1}{m} = \frac{\mathcal{N}(\mathfrak{p})-1}{m}$, we conclude that¹ $(\phi_\alpha, R/\mathfrak{p}) = \alpha^{(\mathcal{N}(\mathfrak{p})-1)/m} \bmod \mathfrak{p}$. This coincides with the definition of the power residue symbol $\left(\frac{\alpha}{\mathfrak{p}}\right)_m$. \square

Example 7.11. Put, again, $K = \mathbb{Q}(\zeta_6, \sqrt[3]{2})$ with order $R = \mathbb{Z}[\zeta_6, \sqrt[3]{2}]$ and the R -ideal $\mathfrak{p}_5 = (5, 3 - \sqrt[3]{2})$, as in Example 7.7. Putting $\alpha = \zeta_6 + 1$, we want to verify that $\left(\frac{\zeta_6+1}{\mathfrak{p}_5}\right)_6 = (\phi_{\zeta_6+1}, R/\mathfrak{p}_5)$, as in Lemma 7.10. The computation of $\left(\frac{\zeta_6+1}{\mathfrak{p}_5}\right)_6$ happens by observing that $\mathcal{N}(\mathfrak{p}_5) = 25$ and computing (using Lemma 7.2)

$$\begin{aligned} (\zeta_6 + 1)^{\frac{\mathcal{N}(\mathfrak{p}_5)-1}{6}} &= (\zeta_6 + 1)^4 = \zeta_6^4 + 4 \cdot \zeta_6^3 + 6 \cdot \zeta_6^2 + 4 \cdot \zeta_6 + 1 \\ &\equiv 9 \cdot \zeta_6 + 9 \equiv -(\zeta_6 + 1) = \zeta_6^5 \bmod \mathfrak{p}_5. \end{aligned}$$

Therefore, $\left(\frac{\alpha}{\mathfrak{p}_5}\right)_6 = \zeta_6^5$. The computation of the signature gives the same result, as can be seen in Figure 7.3. For the computation of the images in that figure; $\phi_{1+\zeta}(\zeta + 2) = (1 + \zeta)(2 + \zeta) = 4 \cdot \zeta + 1 \equiv -\zeta + 1 = -\zeta^2 \pmod{\mathfrak{p}_5}$ and $\phi_{\zeta+1}(\zeta - 2) = (1 + \zeta)(\zeta - 2) = -3 \equiv 2 \pmod{\mathfrak{p}_5}$.

For later purposes, we will need the following lemma, which shows that the signature map $(\cdot, M) : \text{Aut}_{\mathbb{Z}[\zeta_m]}(M) \rightarrow \langle \zeta_m \rangle$ is a group homomorphism.

Lemma 7.12. For two automorphisms ϕ, ψ of an admissible module M , we have

$$(\phi \circ \psi, M) = (\phi, M) \cdot (\psi, M)$$

¹Note that this element $\bar{\alpha}^{(\mathcal{N}(\mathfrak{p})-1)/m}$ coincides with the action of multiplication $x \mapsto \zeta_m^j x$ on R/\mathfrak{p} for some $j \in \mathbb{Z}/m\mathbb{Z}$.

7. The Power Residue Symbol is in ZPP




Figure 7.3.: The computation of the signature $(\phi_{1+\zeta}, R/\mathfrak{p}_5)$ of the map $\phi_{1+\zeta}(x)$, given by the rule $\phi_{1+\zeta}(x) = (1 + \zeta) \cdot x$, as in Example 7.11. By taking the sum of the images' displacements from the chosen representatives (the colored points), we obtain $-1 + 0 + 2 - 2 = -1$. Therefore, we conclude that $(\phi_{1+\zeta}, R/\mathfrak{p}_5) = \zeta_6^{-1} = \zeta_6^5$.

Proof. Choose a representative system S of $M/\langle \zeta_m \rangle$. Then

$$\begin{aligned} (\phi\psi, M) &= \prod_{s \in S} \frac{\phi(\psi(s))}{s^{\phi\psi}} = \prod_{s \in S} \frac{\phi(\psi(s))}{(s^\psi)^\phi} = \prod_{s \in S} \frac{\phi(\psi(s))}{\phi(s^\psi)} \frac{\phi(s^\psi)}{(s^\psi)^\phi} \\ &= \phi \left(\prod_{s \in S} \frac{\psi(s)}{s^\psi} \right) \prod_{s \in S} \frac{\phi(s)}{s^\phi} = \phi((\psi, M)) \cdot (\phi, M) = (\psi, M) \cdot (\phi, M). \end{aligned}$$

□

7.4.4. Invariant Factor Decomposition of R/\mathfrak{b}

Computing the invariant factor decomposition of R/\mathfrak{b} as a module over $\mathbb{Z}[\zeta_m]$ happens by means of the Smith normal form in Dedekind domains (see [Coh99, §1.7]).

This particular Smith normal form algorithm as described in Cohen's book [Coh99, §1.7], needs modules to be represented in terms of *pseudobases*.

Usually, (in a computer algebra system) an R -ideal \mathfrak{b} is represented by means of a basis over \mathbb{Z} instead. We shortly describe here how to obtain such a pseudobasis from a \mathbb{Z} -basis. Let $\mathfrak{b} = \sum_{j=1}^t \mathbb{Z}\beta_j$. Then it is clear that the same set $(\beta_j)_{j \in \{1, \dots, t\}}$ is also a generating set over $\mathbb{Z}[\zeta_m]$, that is: $\mathfrak{b} = \sum_{j=1}^t \mathbb{Z}[\zeta_m]\beta_j$. By using the Hermite normal form over Dedekind domains [Coh99, §1.4] that removes linear dependencies, we arrive at a pseudobasis of \mathfrak{b} over $\mathbb{Z}[\zeta_m]$. The exact same reasoning can be applied to obtain a pseudobasis the ring R as a module over $\mathbb{Z}[\zeta_m]$.

Remark 7.13. *In the undergraduate thesis of Squirrel [Squ97], this step is partially done by computing $\mathbb{Z}[\zeta_m]$ -annihilators of the module R/\mathfrak{b} [Squ97, Ch. 4, §3].*

By [Coh99, §1.7], using a modular Smith normal form, we can deduce that we can find pseudobases for R and \mathfrak{b} of the following shape. $R = \bigoplus_{j=1}^t \mathfrak{s}_j \omega_j$, and $\mathfrak{b} = \bigoplus_{j=1}^t \mathfrak{d}_j \mathfrak{s}_j \omega_j$ where \mathfrak{s}_j are ideals of $\mathbb{Z}[\zeta_m]$, \mathfrak{d}_j are integral ideals of $\mathbb{Z}[\zeta_m]$ satisfying $\mathfrak{d}_{j-1} \subsetneq \mathfrak{d}_j$ for $j \geq 2$ and $\omega_j \in R$. This means that $R/\mathfrak{b} \xrightarrow{\sim} \bigoplus_{j=1}^t \mathbb{Z}[\zeta_m]/\mathfrak{d}_j$.

7.4.5. The Signature is Compatible with Short Exact Sequences

Proposition 7.14. *Let M', M, M'' be admissible $\mathbb{Z}[\zeta_m]$ -modules and let ϕ', ϕ, ϕ'' be $\mathbb{Z}[\zeta_m]$ -module automorphisms of M', M, M'' such that the following diagram commutes.*

$$\begin{array}{ccccccc} 0 & \longrightarrow & M' & \xrightarrow{\iota} & M & \xrightarrow{\pi} & M'' \longrightarrow 0 \\ & & \downarrow \phi' & & \downarrow \phi & & \downarrow \phi'' \\ 0 & \longrightarrow & M' & \xrightarrow{\iota} & M & \xrightarrow{\pi} & M'' \longrightarrow 0 \end{array}$$

Then

$$(\phi', M')(\phi'', M'') = (\phi, M).$$

Proof. Let S' be a representative set of $M'/\langle\zeta_m\rangle$. Extend $\iota(S')$ with a (disjoint) set $S'' \subseteq M$ such that $S = \iota(S') \cup S''$ is a representative set of $M/\langle\zeta_m\rangle$. Then

$$(\phi, M) = \prod_{s \in S} \frac{\phi(s)}{s^\phi} = \prod_{s' \in S'} \frac{\phi(\iota(s'))}{\iota(s')^\phi} \cdot \prod_{s'' \in S''} \frac{\phi(s'')}{(s'')^\phi} = (\phi', M')(\phi'', M''), \quad (7.110)$$

where the last equality is proven in two parts.

(i) As $\phi\iota = \iota\phi'$, we have $\iota(s')^\phi = \iota((s')^{\phi'})$. Therefore,

$$\prod_{s' \in S'} \frac{\phi(\iota(s'))}{\iota(s')^\phi} = \iota\left(\prod_{s' \in S'} \frac{\phi'(s')}{(s')^{\phi'}}\right) = \iota((\phi', M')) = (\phi', M').$$

(ii) Since S'' is distinct from $\iota(S')$, none of the $s'' \in S''$ send to zero under π . Therefore, we can apply π to the rightmost factor in Equation (7.110).

$$\pi\left(\prod_{s'' \in S''} \frac{\phi(s'')}{(s'')^\phi}\right) = \prod_{s'' \in S''} \frac{\pi\phi(s'')}{\pi((s'')^\phi)} = \prod_{s'' \in S''} \frac{\phi(\pi(s''))}{\pi(s'')^{\phi''}} \quad (7.111)$$

As S'' covers all $\langle\zeta_m\rangle$ -orbits of M that do not send to zero under π , the map $S'' \rightarrow \pi(S'')$, $s'' \mapsto \pi(s'')$ is a $|M'|$ -to-one map, i.e., $|\pi(S'')| = |S''|/|M'|$. Also, by surjectivity, $\pi(S'')$ is a representative set for the set $(M'' \setminus 0)/\langle\zeta_m\rangle$. Therefore, Equation (7.111) equals

$$\left(\prod_{t \in \pi(S'')} \frac{\phi''(t)}{t^{\phi''}}\right)^{|M'|} = ((\phi'', M''))^{|M'|} = (\phi'', M''),$$

where the last equality holds because $|M'| \equiv 1$ modulo m and $(\phi'', M'') \in \langle\zeta_m\rangle$. □

Lemma 7.15. *Let R be an order in a number field K with $\mathbb{Z}[\zeta_m] \subseteq R$. Let \mathfrak{b} be an ideal in R , coprime with m . Let $\alpha \in R$ such that $\bar{\alpha} = \alpha \bmod \mathfrak{b} \in (R/\mathfrak{b})^*$ and denote $\phi_\alpha : R/\mathfrak{b} \rightarrow R/\mathfrak{b}$, $x \mapsto \bar{\alpha} \cdot x$. Then*

$$\left(\frac{\alpha}{\mathfrak{b}}\right)_m = (\phi_\alpha, R/\mathfrak{b})$$

Proof. We proceed by induction on the number of different prime ideal factors of \mathfrak{b} . The base case consists of \mathfrak{b} having only one prime divisor, i.e., $\mathfrak{b} = \mathfrak{p}^k$ being a prime power. If $k = 1$, we can apply Lemma 7.10. If $k > 1$, we can construct the following exact sequence

$$0 \rightarrow R/\mathfrak{p} \rightarrow R/\mathfrak{p}^k \rightarrow R/\mathfrak{p}^{k-1} \rightarrow 0$$

where the injection map is defined (non-canonically) by multiplying by an element $\gamma \in \mathfrak{p}^{k-1} \setminus \mathfrak{p}^k$. Then, together with the multiplication-by- α map (which we conveniently write ϕ_α for all rings involved), above exact sequence satisfies the conditions of Proposition 7.14. Therefore, by induction,

$$(\phi_\alpha, R/\mathfrak{p}^k) = (\phi_\alpha, R/\mathfrak{p}^{k-1}) \cdot (\phi_\alpha, R/\mathfrak{p}) = \left(\frac{\alpha}{\mathfrak{p}^{k-1}} \right)_m \left(\frac{\alpha}{\mathfrak{p}} \right)_m = \left(\frac{\alpha}{\mathfrak{p}^k} \right)_m.$$

The induction step consists of \mathfrak{b} being not a prime power. In that case, we write $\mathfrak{b} = \mathfrak{p}^k \mathfrak{c}$ with \mathfrak{p} prime, $k \geq 1$ and $\mathfrak{p} \nmid \mathfrak{c}$, and construct the following exact sequence

$$0 \rightarrow R/\mathfrak{p}^k \rightarrow R/\mathfrak{b} \rightarrow R/\mathfrak{c} \rightarrow 0,$$

where the injection $R/\mathfrak{p}^k \rightarrow R/\mathfrak{b}$ is defined (non-canonically) by multiplying by an element $\gamma \in \mathfrak{c}$ that satisfies $\gamma \equiv 1$ modulo \mathfrak{p}^k . Again denoting ϕ_α for multiplication by α in all of the rings involved, this exact sequence satisfies the conditions of Proposition 7.14. Therefore, by induction,

$$(\phi_\alpha, R/\mathfrak{b}) = (\phi_\alpha, R/\mathfrak{p}^k) \cdot (\phi_\alpha, R/\mathfrak{c}) = \left(\frac{\alpha}{\mathfrak{p}^k} \right)_m \cdot \left(\frac{\alpha}{\mathfrak{c}} \right)_m = \left(\frac{\alpha}{\mathfrak{b}} \right)_m.$$

□

7.4.6. The Determinant Formula

Let $M = (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ for some ideal \mathfrak{c} of $\mathbb{Z}[\zeta_m]$ and some $t \in \mathbb{N}_{>0}$. Then any automorphism $\phi : M \rightarrow M$ can be described as a non-degenerate $t \times t$ matrix with coefficients in $\mathbb{Z}[\zeta_m]/\mathfrak{c}$, which we call M_ϕ .

Lemma 7.16. *We have*

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = \left(\frac{\det(M_\phi)}{\mathfrak{c}} \right)_{m, \mathbb{Q}(\zeta_m)}$$

Proof. We prove the statement first for $\mathfrak{c} = \mathfrak{p}$ a prime ideal. In that case the matrix M_ϕ has coefficients in the field $\mathbb{Z}[\zeta_m]/\mathfrak{p}$. Matrices over fields can be decomposed into $M_\phi = ULU'$, where U, U' are upper triangular and L is lower triangular, by means of Gaussian elimination. We denote ϕ_U, ϕ_L, ϕ'_U for their associated maps on $(\mathbb{Z}[\zeta_m]/\mathfrak{p})^t$. We have the exact sequence

$$0 \rightarrow \mathbb{Z}[\zeta_m]/\mathfrak{p} \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^{t-1} \rightarrow 0$$

where the injection map is just $x \mapsto (x, 0, \dots, 0)$ and the projection map projects on the last $t-1$ coordinates. By the (upper/lower) triangular shape of the matrix U of ϕ_U and by induction, one can deduce that

$$(\phi_U, M) = \left(\frac{\det(U)}{\mathfrak{p}} \right),$$

and the same for U' and L . Therefore,

$$\begin{aligned} (\phi, M) &= (\phi_U \phi_L \phi'_U, M) = (\phi_U, M)(\phi_L, M)(\phi'_U, M) \\ &= \left(\frac{\det(U)}{\mathfrak{p}} \right) \left(\frac{\det(L)}{\mathfrak{p}} \right) \left(\frac{\det(U')}{\mathfrak{p}} \right) = \left(\frac{\det(ULU')}{\mathfrak{p}} \right) = \left(\frac{M_\phi}{\mathfrak{p}} \right). \end{aligned}$$

This proves the statement for \mathfrak{c} being a prime ideal. For the general case, write $\mathfrak{c} = \mathfrak{p}\mathfrak{a}$, and construct the exact sequence

$$0 \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow 0 \quad (7.112)$$

where the injection map is defined by scalar multiplication by $\varpi \in \mathfrak{p} \setminus \mathfrak{p}^2$ and the projection map just takes the entries modulo \mathfrak{p} .

Let $\phi' : (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t$ respectively $\phi'' : (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t$ be the map defined by reducing the entries of the matrix $M_\phi \in (\mathbb{Z}[\zeta_m]/\mathfrak{c})^{t \times t}$ modulo \mathfrak{a} respectively \mathfrak{p} . Then Equation (7.112) satisfies the requirements of Proposition 7.14, therefore

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = (\phi', (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t) \cdot (\phi'', (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t)$$

$$= \left(\frac{\det(M_\phi)}{\mathfrak{a}} \right) \left(\frac{\det(M_\phi)}{\mathfrak{p}} \right) = \left(\frac{\det(M_\phi)}{\mathfrak{c}} \right).$$

Here, we used the induction hypothesis, the fact that

$$\det(M_{\phi'}) = \det(M_\phi \bmod \mathfrak{a}) = \det(M_\phi) \bmod \mathfrak{a},$$

and the similar statement for \mathfrak{p} . □

7.4.7. Applying Induction on the Components of the Module

Lemma 7.17. *Let $R \subseteq K$ be a number ring containing a primitive m -th root of unity ζ_m and let $\mathfrak{b} \subseteq R$ be an ideal coprime with m . Let*

$$R/\mathfrak{b} = \gamma_1 \mathbb{Z}[\zeta_m]/\mathfrak{d}_1 \oplus \cdots \oplus \gamma_k \mathbb{Z}[\zeta_m]/\mathfrak{d}_k, \quad (7.113)$$

be the invariant factor decomposition of R/\mathfrak{b} with $\mathfrak{d}_j = \prod_{\ell \leq j} \mathfrak{c}_\ell$. Then we have, for all $\alpha \in R$ coprime with both \mathfrak{b} and m ,

$$\left(\frac{\alpha}{\mathfrak{b}} \right)_{m,K} = \prod_{j=1}^k \left(\frac{d_j}{\mathfrak{c}_j} \right)_{m, \mathbb{Q}(\zeta_m)},$$

where $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are specific determinants of $k-j+1 \times k-j+1$ matrices with coefficients in $\mathbb{Z}[\zeta_m]/\mathfrak{c}_j$.

Proof. Denoting $M = R/\mathfrak{b}$, we have the following exact sequence

$$0 \rightarrow M/(\mathfrak{c}_1 M) \rightarrow M \rightarrow \mathfrak{c}_1 M \rightarrow 0,$$

where \mathfrak{c}_1 is the first factor in the invariant factor decomposition (Equation (7.107)). Because of the compatibility of the signature with short exact sequences, it is enough to compute the signatures $(m_\alpha, M/(\mathfrak{c}_1 M))$ and $(m_\alpha, \mathfrak{c}_1 M)$.

The first module, $M/(\mathfrak{c}_1 M)$, can be shown to be isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c}_1)^k$, and therefore the determinant formula applies (see Equation (7.108)).

The last module, $\mathfrak{c}_1 M$, can be shown to have less ‘components’ than M itself; $k - 1$ instead of k .

$$\mathfrak{c}_1 M = \bigoplus_{j=1}^{k-1} \gamma_j \mathbb{Z}[\zeta_m] / \tilde{\mathfrak{d}}_j,$$

where $\tilde{\mathfrak{d}}_j = \mathfrak{d}_j / \mathfrak{c}_1$, and where \mathfrak{d}_j are obtained from the invariant factor decomposition of $M = R/\mathfrak{b}$. □

7.4.8. Conclusion

By induction, we can therefore conclude that the computation of

$$\left(\frac{\alpha}{\mathfrak{b}} \right)_{m,R} = (\phi_\alpha, R/\mathfrak{b})$$

reduces to the computation of k power residue symbols $\left(\frac{d_j}{\mathfrak{c}_j} \right)_{m, \mathbb{Q}(\zeta_m)}$ (for $j \in \{1, \dots, k\}$) in the cyclotomic field $\mathbb{Q}(\zeta_m)$. Here, \mathfrak{c}_j are the invariant factors of the module R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module as in Equation (7.113) and $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are determinants of associated automorphisms. We thus proved the following statement.

Theorem 7.18 (Lenstra, Squirrel). *Let $R \subseteq K$ be a number ring of a number field containing the m -th root of unity ζ_m . Let $\mathfrak{b} \subseteq R$ be an ideal coprime with m and let $\alpha \in R$ be an element of coprime with \mathfrak{b} and m . Then the computation of the power residue symbol $\left(\frac{\alpha}{\mathfrak{b}} \right)_{m,R}$ reduces to at most $\log(\mathcal{N}(\mathfrak{b}))$ computations of the power residue symbols $\left(\frac{d_j}{\mathfrak{c}_j} \right)$ in $\mathbb{Q}(\zeta_m)$, where the d_j and \mathfrak{c}_j are bounded in size by the size of \mathfrak{b} and α .*

7.5. Computing the Power Residue Symbol in Cyclotomic Fields

7.5.1. Main Idea

Before explaining an algorithm in full detail, it is often insightful to give a simplified version first. The simplified version of the algorithm that computes power residue symbols $\left(\frac{\alpha}{\mathfrak{b}}\right)$ with an element $\alpha \in \mathbb{Z}[\zeta_m]$ and an integral ideal \mathfrak{b} of $\mathbb{Z}[\zeta_m]$ essentially proceeds by two steps. An essential part of the algorithm is the idea that prime ideals ‘occur quite often’ in cyclotomic fields. This is a consequence of the density of primes of norm N being around $\frac{1}{\rho_K \log N}$ and the fact that the residue ρ_K of the Dedekind zeta function of cyclotomic fields at $s = 1$ is polynomially bounded (see Appendix A.2).

Step 1: Reducing the symbol $\left(\frac{\alpha}{\mathfrak{b}}\right)$ to a ‘principal’ symbol $\left(\frac{\alpha}{\beta}\right)$.

This happens by repeatedly sampling random $\beta \in \mathfrak{b}$ until the ideal $(\beta)/\mathfrak{b}$ is equal to some prime ideal \mathfrak{p} of $\mathbb{Z}[\zeta_m]$. In that case, write $(\beta) = \mathfrak{p}\mathfrak{b}$ and use the multiplicative property of the power residue symbol to obtain $\left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\mathfrak{b}}\right) \cdot \left(\frac{\alpha}{\mathfrak{p}}\right)$. By the fact that there exists an efficiently computable formula (see Lemma 7.2) for power residue symbols with a prime ideal as the bottom input, the symbol $\left(\frac{\alpha}{\mathfrak{p}}\right)$ is efficiently computable.

Therefore, provided that such a suitable $\beta \in \mathfrak{b}$ can be efficiently found, the above procedure reduces the computation of the symbol $\left(\frac{\alpha}{\mathfrak{b}}\right)$ to the computation of a power residue symbol $\left(\frac{\alpha}{\beta}\right)$ where the bottom input β is an *element* in $\mathbb{Z}[\zeta_m]$ instead of a generic ideal.

Step 2: Evaluating the symbol $\left(\frac{\alpha}{\beta}\right)$ by shifting β .

This happens by sampling random $\kappa \in \mathbb{Z}[\zeta_m]$ until the shifted element $\beta + \kappa m^m \alpha = \varpi$ is a *prime element*. As the power residue symbol $\left(\frac{\alpha}{\beta}\right)$ with

$\alpha, \beta \in \mathbb{Z}[\zeta_m]$ satisfies the ‘shifting property’, (see Lemma 7.3) we have

$$\left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\beta + \kappa m^m \alpha}\right) = \left(\frac{\alpha}{\varpi}\right).$$

Because (ϖ) is a prime ideal, there exists an efficiently computable formula for the symbol $(\frac{\alpha}{\varpi})$ (see Lemma 7.2). Therefore, $(\frac{\alpha}{\beta})$ can also be computed efficiently, provided that one indeed can find a $\kappa \in \mathbb{Z}[\zeta_m]$ such that $\beta + \kappa m^m \alpha$ is a prime element in $\mathbb{Z}[\zeta_m]$.

Discussion

It is clear that the first step only works whenever sampling a random $\beta \in \mathfrak{b}$ results sufficiently often in an ideal $(\beta)/\mathfrak{b}$ that is prime. In other words, the probability that $(\beta)/\mathfrak{b}$ is prime should be high enough. Likewise, the second step only works whenever sampling random $\kappa \in \mathbb{Z}[\zeta_m]$ results sufficiently often in an element $\beta + \kappa m^m \alpha$ that is prime.

It turns out to be notoriously hard to estimate these probabilities whenever \mathfrak{b} and β are *fixed*. However, if both \mathfrak{b} and β are appropriately *random* instead, one can actually lower bound these probabilities by means of *Landau’s prime ideal theorem*. This theorem can be informally expressed by saying that there are many prime ideals among the ideals in $\mathbb{Z}[\zeta_m]$. In other words, if one takes a ‘random ideal’ in $\mathbb{Z}[\zeta_m]$, there is a reasonable probability that it is a prime ideal.

So, in order to be fully able to estimate the success probability of the algorithm, we will need to appropriately *randomize* the lower input of the power residue symbol. With this adequate randomization, which will be done by means of a random walk as in Chapter 4 (thus relying on the Extended Riemann Hypothesis), one obtains the provable, full algorithm.

Remark 7.19. *In an actual implementation, one should not use this chapter’s provable algorithm. Instead, one should use the heuristic variant of it described in [BP17; Boe16]. A specific blend between the provable and the heuristic variant that uses Artin reciprocity (see Lemma 7.3) might also be*

Algorithm 8: POWERRESIDUESYMBOL(α, \mathfrak{b}, m), the computation of the symbol $(\frac{\alpha}{\mathfrak{b}})$

Require:

- An integer $m > 1$ defining the cyclotomic field $\mathbb{Q}(\zeta_m)$ of degree n .
- An integral element $\alpha \in \mathbb{Z}[\zeta_m]$ coprime with m .
- An integral ideal $\mathfrak{b} \subseteq \mathbb{Z}[\zeta_m]$ coprime with α and m ,

Ensure: $(\frac{\alpha}{\mathfrak{b}}) \in \langle \zeta_m \rangle$, or failure.

- 1: Put $\mathfrak{m} = m^m \cdot (\alpha)$ as the modulus.
- 2: Apply the sampling Algorithm 7 with \mathfrak{b} , \mathfrak{m} , $\tau = 1$ and $1/\varepsilon = \max(2^n, n^{5+1}(n + \log |\mathcal{N}(\alpha)|))$ to sample an element $\beta \in \tilde{\mathfrak{b}} \cap (1 + \mathfrak{m})$, where $\tilde{\mathfrak{b}} = \mathfrak{b} \prod_j \mathfrak{p}_j$ comes from the sampling algorithm.
- 3: **return** $(\frac{\alpha}{\mathfrak{p}})^{-1} \cdot \prod_j (\frac{\alpha}{\mathfrak{p}_j})^{-1}$ if $\beta/\tilde{\mathfrak{b}} = \mathfrak{p}$ is prime, using the formula for the power residue symbol above prime ideals (Lemma 7.2).
- 4: **return** *failure* otherwise.

interesting to implement, because it avoids the need for the computation of Hilbert symbols. Such an implementation (that relies on Artin reciprocity and not Hilbert reciprocity) might therefore even be used to compute Hilbert symbol due to a ‘global-to-local’ principle (see also Section 7.6.1).

7.5.2. The Full Algorithm

Lemma 7.20 (ERH). *Assuming the Riemann Hypothesis for Hecke L -functions on cyclotomic fields, Algorithm 8 is correct and runs in time polynomial in $m, \log |\mathcal{N}(\alpha)|$ and $\log \mathcal{N}(\mathfrak{b})$. Furthermore, Algorithm 8 has success probability at least*

$$\Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log |\mathcal{N}(\alpha)|)}\right)$$

Proof. We start with proving the correctness of Algorithm 8, i.e., that the algorithm computes the symbol $(\frac{\alpha}{\mathfrak{b}})$ if it does not fail. This is proven by the sequence of equalities in Equation (7.114), which uses the multiplicative property of the power residue symbol (see Lemma 7.4) and the fact that the power residue symbol is trivial on the ray $K^{\mathfrak{m},1}$ with $\mathfrak{m} = m^m(\alpha)$ (see Lemma 7.3). So, since $\beta \in K^{\mathfrak{m},1}$ (i.e., $(\frac{\alpha}{\beta}) = 1$) and $(\beta) = \mathfrak{p}\tilde{\mathfrak{b}} = \mathfrak{p}\mathfrak{b} \prod_j \mathfrak{p}_j$, one obtains

$$1 = \left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \left(\frac{\alpha}{\tilde{\mathfrak{b}}}\right) = \left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \left(\frac{\alpha}{\mathfrak{b}}\right) \cdot \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right). \quad (7.114)$$

The correctness of the algorithm follows by rearranging terms to get an expression for $(\frac{\alpha}{\mathfrak{b}})$.

For the success probability, we need to estimate the probability that $(\beta)/\tilde{\mathfrak{b}}$ is a prime ideal in step 3. By the correspondence theorem between sampling probability and ideal density (see Theorem 6.21) we know that the probability of $(\beta)/\tilde{\mathfrak{b}}$ being prime equals at least $\delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon$, where $\mathcal{S}^{\mathfrak{m}} = \{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \mid \mathfrak{p} \text{ prime}\}$. By Lemma 7.5, the fact that $r^n \geq \mathcal{N}(\mathfrak{m}) \geq 16 \cdot \omega(\mathfrak{m})^2$, Writing out the instantiation for r in Algorithm 7, using $|\Delta_K|^{3/(2n)} \leq n^{3/2}$ for cyclotomic fields K , we have

$$\begin{aligned} r &= 4 \cdot 2^n \cdot n^{3/2} \cdot |\Delta_K|^{\frac{3}{2n}} \cdot \mathcal{N}(\mathfrak{m})^{1/n} \leq 4 \cdot 2^n \cdot n^3 \cdot \mathcal{N}(\mathfrak{m})^{1/n} \\ &\leq 2^{n+2} \cdot n^3 \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathcal{N}(\alpha)^{1/n}, \end{aligned}$$

I.e., $\log(r^n) \leq n(n+2)\log(2) + 3n\log(n) + n^2\log n + \log|\mathcal{N}(\alpha)| = O(n^2\log n + \log|\mathcal{N}(\alpha)|)$. Then, we have that the success probability is lower bounded (see Theorem 6.21) by

$$\delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon \geq \frac{1}{3 \cdot \rho_K \cdot n \cdot \log r} - \varepsilon \geq \frac{1}{\rho_K \cdot n \cdot (n^2 \log n + \log|\mathcal{N}(\alpha)|)} - \varepsilon$$

We show in Appendix A.2 that $\rho_K = O(n^4)$ (the hidden constant is $e^{15} \approx 3.3 \cdot 10^6$). By the instantiation $1/\varepsilon = \max(2^n, n^{5+1}(n^2 \log n + \log|\mathcal{N}(\alpha)|))$ we then have,

$$\begin{aligned} \delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon &= \Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log|\mathcal{N}(\alpha)|)}\right) - \varepsilon \\ &= \Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log|\mathcal{N}(\alpha)|)}\right) \end{aligned}$$

As Algorithm 7 is polynomial in its input size and in $\log(1/\varepsilon)$, it is enough to show that $\log \mathcal{N}(\mathfrak{p})$, $\log \mathcal{N}(\mathfrak{p}_j)$ and $\log(1/\varepsilon)$ are polynomially bounded in m , $\log |\mathcal{N}(\alpha)|$ and $\log \mathcal{N}(\mathfrak{b})$, in order to prove that Algorithm 8 runs in polynomial time.

Note that $\mathfrak{m} = m^m(\alpha)$, therefore $\log(\mathcal{N}(\mathfrak{m})) = \text{poly}(n, \log |\mathcal{N}(\alpha)|)$ is polynomially bounded. The logarithm of the inverse error $\log(1/\varepsilon)$ is easily shown to be polynomially bounded as well. Also $N, \log B$ and $\log r$ from Algorithm 7 with the instantiation of ε are polynomially bounded by $m, \log |\Delta_K| = O(m), \log(1/\varepsilon)$ and $\log \mathcal{N}(\mathfrak{d})$. So $\log \mathcal{N}(\mathfrak{p}_j) \leq \log B$ are polynomially bounded.

The largest prime, \mathfrak{p} , satisfies $\log \mathcal{N}(\mathfrak{p}) \leq \log(|\mathcal{N}(\beta)| / \mathcal{N}(\mathfrak{b})) \leq N \log B + n \log r$, by Algorithm 7. Therefore, all relevant quantities are polynomially bounded, thus the entire algorithm runs within polynomial time. \square

Theorem 7.21. *Let $K \supseteq \mathbb{Q}(\zeta_m)$ be a number field and let $R \subseteq K$ be an order in that number field. Assume the Extended Riemann Hypothesis for Hecke-L functions of the cyclotomic number field $\mathbb{Q}(\zeta_m)$.*

Then, there exists an algorithm that computes the power residue symbol $(\frac{\alpha}{\mathfrak{b}})$ for all elements $\alpha \in R$ and ideals $\mathfrak{b} \subseteq R$, within time polynomial in $\log |\Delta_K|$, $[K : \mathbb{Q}]$, $\text{size}(\alpha)$ and $\text{size}(\mathfrak{b})$.

Proof. Follows immediately from Lemma 7.20 and the reduction from Lenstra and Squirrel (Theorem 7.18). \square

7.6. Discussion

7.6.1. Computing Hilbert Symbols Using Power Residue Symbols

Because the algorithm in this chapter does not use the computation of Hilbert symbols (as opposed to the heuristic algorithm in [BP17; Boe16]), one can reverse the roles and use the computation of power residue symbols

to derive information about the associated Hilbert symbols in a number field K containing $\mathbb{Q}(\zeta_m)$ in the following way [Neu85, Ch. IV, §9].

$$\prod_{\mathfrak{p}|m\infty} (\alpha, \beta)_{\mathfrak{p}} = \left(\frac{\alpha}{\beta} \right)_m \left(\frac{\beta}{\alpha} \right)_m^{-1}.$$

To compute $(\alpha, \beta)_{\mathfrak{q}}$ for a fixed chosen $\mathfrak{q} \mid m$, one picks, using the Chinese remainder theorem, an element $\gamma \in \mathcal{O}_K$ that satisfies $\gamma \equiv 1$ modulo \mathfrak{p}^{d^2} for $\mathfrak{p} \mid m$ and $\mathfrak{p} \neq \mathfrak{q}$, and $\gamma \equiv \beta$ modulo \mathfrak{q}^{d^2} , where $d = [K : \mathbb{Q}]$ is the degree of the number field K . In that case, $(\alpha, \gamma)_{\mathfrak{p}} = 1$ for $\mathfrak{p} \neq \mathfrak{q}$ and $(\alpha, \gamma)_{\mathfrak{q}} = (\alpha, \beta)_{\mathfrak{q}}$, and therefore

$$(\alpha, \beta)_{\mathfrak{q}} = (\alpha, \gamma)_{\mathfrak{q}} = \prod_{\mathfrak{p}|m\infty} (\alpha, \gamma)_{\mathfrak{p}} = \left(\frac{\alpha}{\gamma} \right)_m \left(\frac{\gamma}{\alpha} \right)_m^{-1}.$$

In above reasoning, we use the following lemma.

Lemma 7.22. *Let $K_{\mathfrak{p}}$ be the completion of a number field $K \supseteq \mathbb{Q}(\zeta_m)$ of degree $d = [K : \mathbb{Q}]$ with respect to the finite prime $\mathfrak{p} \mid m$, and let $(\cdot, \cdot)_{\mathfrak{p}} : K_{\mathfrak{p}}^* \times K_{\mathfrak{p}}^* \rightarrow \langle \zeta_m \rangle$ denote the Hilbert symbol on this completion. Then*

$$(\alpha, 1 + \pi^{d^2})_{\mathfrak{p}} = 1 \text{ for all } \pi \in \mathfrak{p}.$$

Proof. As $(\alpha, \cdot)_{\mathfrak{p}} : K_{\mathfrak{p}} \rightarrow \langle \zeta_m \rangle$ equals the Artin symbol (or norm residue symbol) of the extension $K_{\mathfrak{p}}(\sqrt[d]{\alpha}) : K_{\mathfrak{p}}$ [Neu85, Ch. 3, Prop. 5.1], it suffices to show that $1 + \pi^{d^2} \in \mathcal{N}_{K_{\mathfrak{p}}(\sqrt[d]{\alpha})/K_{\mathfrak{p}}}(K_{\mathfrak{p}}(\sqrt[d]{\alpha}))$ for all $\pi \in \mathfrak{p}$ [Neu85, Ch. 3, Prop. 5.2iii]. In other words, we need to show that the conductor $\mathfrak{f}_{K_{\mathfrak{p}}(\sqrt[d]{\alpha})/K_{\mathfrak{p}}}$ of this local Kummer extension divides \mathfrak{p}^{d^2} . By using local computations and Hensel's lemma [CS08, Eq. (3.11)], we know that

$$\text{ord}_{\mathfrak{p}}(\mathfrak{f}_{K_{\mathfrak{p}}(\sqrt[d]{\alpha})/K_{\mathfrak{p}}}) \leq d(1 + \log(d/m)) + 1 \leq d^2.$$

□

This leads to the following corollary.

Corollary 7.23. *Assuming the Extended Riemann Hypothesis for Hecke- L functions on cyclotomic fields, Hilbert symbols can be computed within time polynomial in the input size.*

This corollary is quite weak compared to the much stronger results of Bouw [Bou21]; his algorithm is unconditional (i.e., does not require any variant of the Riemann Hypothesis), deterministic, and his algorithm's studied complexity is way more explicit. Though no real-life comparison has been made yet, I suspect Bouw's algorithm to run significantly faster than the method described above.

7.6.2. Computing Artin Symbols in the Same Fashion

A very similar algorithm as Algorithm 8 could in principle be used to compute Artin symbols $\left(\frac{\cdot}{L/K}\right)$ for abelian extensions L/K . The main caveat is that the residue ρ_K of the Dedekind zeta function $\zeta_K(s)$ of K at $s = 1$ might be too large, i.e., not polynomially bounded. Such a large residue would make such an algorithm not feasible, as the success probability depends inversely on this residue ρ_K .

For the sake of completeness, we do spell out a proposal for an algorithm computing Artin symbols in Algorithm 9. We would like to stress that no guarantee on the running time is given, except maybe whenever the residue ρ_K is polynomially bounded. In that case, the proof resembles that of Lemma 7.20.

Remark 7.24. *To compute the Frobenius element $\left(\frac{\mathfrak{p}}{L/K}\right) \in G = \text{Gal}(L/K)$ for a prime \mathfrak{p} as in Line 3 of Algorithm 9, one goes through the following lines.*

- Compute $\mathfrak{P} \subseteq \mathcal{O}_L$, any prime ideal above $\mathfrak{p} \subseteq \mathcal{O}_K$.
- Compute a primitive element $\alpha \in L$, i.e., an $\alpha \in L$ such that $L = K(\alpha)$, by means of linear algebra.
- Compute $\alpha^q \bmod \mathfrak{P}$, where $q = |\mathcal{O}_K/\mathfrak{p}|$.
- Output a $g \in G = \text{Gal}(L/K)$ for which holds $\alpha^q \equiv g(\alpha) \bmod \mathfrak{P}$.

Algorithm 9: $\text{ArtinSymbol}(\mathfrak{b}, L, K)$, the computation of the Artin symbol $\left(\frac{\mathfrak{b}}{L/K}\right) \in \text{Gal}(L/K)$

Require:

- A number field extension L/K , where both L and K are defined by a defining polynomial over \mathbb{Q} .
- An integral ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ coprime with Δ_L .
- For all $g \in G = \text{Gal}(L/K)$ and $\alpha \in L$, an efficient algorithm that computes $g(\alpha) \in L$.

Ensure: $\left(\frac{\alpha}{\mathfrak{b}}\right) \in \langle \zeta_m \rangle$, or failure.

- 1: Put $\mathfrak{m} = \Delta_{L/K}$ the relative discriminant of the extension L/K as the modulus.
- 2: Apply the sampling Algorithm 7 with \mathfrak{b} , \mathfrak{m} , $\tau = 1$ and $1/\varepsilon = \max(2^n, \rho_K \cdot n \cdot (n^2 \log n + n \log \mathcal{N}(\mathfrak{m})))$ to sample an element $\beta \in \tilde{\mathfrak{b}} \cap (1 + \mathfrak{m})$, where $\tilde{\mathfrak{b}} = \mathfrak{b} \prod_j \mathfrak{p}_j$ comes from the sampling algorithm.
- 3: **return** $\left(\frac{\mathfrak{p}}{L/K}\right)^{-1} \cdot \prod_j \left(\frac{\mathfrak{p}_j}{L/K}\right)^{-1}$ if $\beta/\tilde{\mathfrak{b}} = \mathfrak{p}$ is prime, using the formula for the Artin symbol for prime ideals (‘Frobenius element’, see [Neu85, Ch. IV, §8]).
- 4: **return** *failure* otherwise.

Remark 7.25. *The approach of Algorithm 9 is not expected to work for number fields with large Dedekind residue ρ_K . Though, we might enlarge the set of ‘good’ ideals \mathcal{S} by also including ‘near primes’, which are ideals that are a product of a large prime ideal and several smaller prime ideals; in other words, a large prime ideal times a smooth ideal.*

This might increase the local density of \mathcal{S} significantly in some cases, maybe even to the point that the Algorithm 9 succeeds within polynomial time even though ρ_K is not small.

An open question arising here is: What exactly does a large residue ρ_K mean? If it just implies more frequent small primes or more (higher) prime powers, it does not affect the Artin symbol algorithm. If it, on the other hand, implies a scarcity of easy-to-factor ideals, it does affect the Artin symbol algorithm. Are there means to distinguish these two cases?

