Chapter 6

Algorithms for Algebraic Number Theory II

We now leave the realm of quadratic fields where the main computational tasks of algebraic number theory mentioned at the end of Chapter 4 were relatively simple (although as we have seen many conjectures remain), and move on to general number fields.

We first discuss practical algorithms for computing an integral basis and for the decomposition of primes in a number field K, essentially following a paper of Buchmann and Lenstra [Buc-Len], except that we avoid the explicit use of Artinian rings. We then discuss algorithms for computing Galois groups (up to degree 7, but see also Exercise 15). As examples of number fields of higher degree we then treat cyclic and pure cubic fields. Finally, in the last section of this chapter, we give a complete algorithm for class group and regulator computation which is sufficient for dealing with fields having discriminants of reasonable size. This algorithm also gives a system of fundamental units if desired.

6.1 Computing the Maximal Order

Let $K = \mathbb{Q}[\theta]$ be a number field, where θ is a root of a monic polynomial $T(X) \in \mathbb{Z}[X]$. Recall that $\mathbb{Z}_K$ has been defined as the set of algebraic integers belonging to K, and that it is called the maximal order since it is an order in K containing every order of K. We will build it up by starting from a known order (in fact from $\mathbb{Z}[\theta]$) and by successively enlarging it.

6.1.1 The Pohst-Zassenhaus Theorem

The main tool that we will use for enlarging an order is the Pohst-Zassenhaus Theorem 6.1.3 below. We first need a few basic results and definitions.

Definition 6.1.1. *Let $\mathcal{O}$ be an order in a number field K and let p be a prime number.*

(1) *We will say that $\mathcal{O}$ is* p-maximal *if $[\mathbb{Z}_K : \mathcal{O}]$ is not divisible by p.*

(2) *We define the* p-radical *I_p as follows.*

$$I_p = \{x \in \mathcal{O} \mid \exists m \geq 1 \text{ such that } x^m \in p\mathcal{O}\}$$

Proposition 6.1.2. *Let $\mathcal{O}$ be an order in a number field K and let p be a prime number.*

(1) *The p-radical I_p is an ideal of $\mathcal{O}$.*
(2) *We have*

$$I_p = \prod_{1 \le i \le g} \mathfrak{p}_i$$

the product being over all distinct prime ideals $\mathfrak{p}_i$ of $\mathcal{O}$ which lie above p.
(3) *There exists an integer m such that $I_p^m \subset p\mathcal{O}$.*

Proof. For (1), the only thing which is not completely trivial is that I_p is stable under addition. If $x^m \in p\mathcal{O}$ and $y^n \in p\mathcal{O}$, then clearly $(x+y)^{n+m} \in p\mathcal{O}$ as we see by using the binomial theorem.

For (2) note that since $\mathfrak{p}_i$ lies above p then $p\mathcal{O} \subset \mathfrak{p}_i$. So, if $x \in I_p$ there exists an m such that $x^m \in p\mathcal{O} \subset \mathfrak{p}_i$, and hence $x \in \mathfrak{p}_i$ by definition of a prime ideal. By Proposition 4.6.4 this shows that $x \in \bigcap_{1 \le i \le g} \mathfrak{p}_i = \prod_{1 \le i \le g} \mathfrak{p}_i$ since the distinct maximal ideals $\mathfrak{p}_i$ are pairwise coprime.

Conversely, assume that $x \in \prod_{1 \le i \le g} \mathfrak{p}_i$. By definition, the set of ideals of $\mathcal{O}$ containing $p\mathcal{O}$ is in canonical one-to-one correspondence with the ideals of the finite quotient ring $R = \mathcal{O}/p\mathcal{O}$. We will use this at length later. For now, note that it implies that this set is finite, and in particular the ideals $\alpha^n R$ are finite in number, where α is the class of x in R. In particular, there exists an n such that $\alpha^n R = \alpha^{n+1} R$, i.e. $\alpha^n(1 - \alpha\beta) = 0$ for some $\beta \in R$. By assumption, α belongs to all the maximal ideals $\overline{\mathfrak{p}_i}$ of R hence $(1 - \alpha\beta)$ cannot belong to any of them, otherwise 1 would also, which is impossible. It follows that the ideal $(1 - \alpha\beta)R$, not being contained in any maximal ideal, must be equal to R, i.e. $1 - \alpha\beta$ is invertible R. The equality $\alpha^n(1 - \alpha\beta) = 0$ thus implies that $\alpha^n = 0$ in R, i.e. that $x^n \in p\mathcal{O}$ or again that $x \in I_p$ as was to be proved.

Finally, for (3) note that since I_p is an ideal of an order in a number field it has a finite $\mathbb{Z}$-basis x_i for $1 \le i \le n$. For each x_i there exists an m_i such that $x_i^{m_i} \in p\mathcal{O}$, and if we set $m = \sum_{1 \le i \le n} m_i$ it is clear that $I_p^m \subset p\mathcal{O}$, again by the binomial theorem. □

The procedure that we will use to obtain the maximal order is to start with $\mathcal{O} = \mathbb{Z}[\theta]$ and enlarge it for successive primes so as to get an order which is p-maximal for every p, hence which will be the maximal order. The enlarging procedure which we will use, due to Pohst and Zassenhaus, is based on the following theorem.

Theorem 6.1.3. *Let $\mathcal{O}$ be an order in a number field K and let p be a prime number. Set*

$$\mathcal{O}' = \{x \in K | x I_p \subset I_p\} .$$

Then either $\mathcal{O}' = \mathcal{O}$, in which case $\mathcal{O}$ is p-maximal, or $\mathcal{O}' \supsetneq \mathcal{O}$ and $p \mid [\mathcal{O}' : \mathcal{O}] \mid p^n$.

Proof. Since I_p is an ideal, it is clear that $\mathcal{O}'$ is a ring containing $\mathcal{O}$. Furthermore, since $p \in I_p$, $x \in \mathcal{O}'$ implies that $xp \in I_p \subset \mathcal{O}$ and hence $\mathcal{O} \subset \mathcal{O}' \subset \frac{1}{p}\mathcal{O}$. This shows that $\mathcal{O}'$ has maximal rank, i.e. is an order in K, and it also shows that $[\mathcal{O}' : \mathcal{O}] | p^n$.

We now assume that $\mathcal{O}' = \mathcal{O}$. Define

$$\mathcal{O}_p = \{x \in \mathbb{Z}_K | \exists j \ge 1, p^j x \in \mathcal{O}\} .$$

It is clear that $\mathcal{O} \subset \mathcal{O}_p$ and that $\mathcal{O}_p$ is an order. Furthermore, $\mathcal{O}_p$ is p-maximal. Indeed, if p divides the index $[\mathbb{Z}_K : \mathcal{O}_p]$, then there exists $x \in \mathbb{Z}_K$ such that $x \notin \mathcal{O}_p$ but $px \in \mathcal{O}_p$. The definition of $\mathcal{O}_p$ shows that this is impossible.

We are now going to show that $\mathcal{O}_p = \mathcal{O}$. Since $\mathcal{O}_p$ is an order, it is finitely generated over $\mathbb{Z}$. Hence there exists an $r \ge 1$ such that $p^r \mathcal{O}_p \subset \mathcal{O}$ (take r to be the maximum of the j such that $p^j x_i \in \mathcal{O}$ for a finite generating set (x_i) of $\mathcal{O}_p$). Since $I_p^m \subset p\mathcal{O}$ it follows that $\mathcal{O}_p I_p^{mr} \subset \mathcal{O}$. Assume by contradiction that $\mathcal{O}_p \neq \mathcal{O}$, hence $\mathcal{O}_p \not\subset \mathcal{O}$. Let n be the largest index such that $\mathcal{O}_p I_p^n \not\subset \mathcal{O}$ (hence n exists and $0 \le n < mr$). We thus also have $\mathcal{O}_p I_p^{n+1} \subset \mathcal{O}$. Choose any $x \in \mathcal{O}_p I_p^n \setminus \mathcal{O}$. Then $x I_p \subset \mathcal{O}$. Since $\mathcal{O}_p I_p^{n+m+1} \subset I_p^m \subset p\mathcal{O}$ it follows that if $y \in I_p$, then $(xy)^{n+m+1} \in p\mathcal{O}$ hence that $xy \in I_p$, so $x I_p \subset I_p$ thus showing that $x \in \mathcal{O}'$. This is a contradiction since $x \notin \mathcal{O}$ and we have assumed that $\mathcal{O}' = \mathcal{O}$. This finishes the proof of Theorem 6.1.3. □

(I thank D. Bernardi for the final part of the proof.)

6.1.2 The Dedekind Criterion

From the Pohst-Zassenhaus theorem, starting from a number field $K = \mathbb{Q}(\theta)$ defined by a monic polynomial $T \in \mathbb{Z}[X]$, we will enlarge the order $\mathbb{Z}[\theta]$ for every prime p such that p^2 divides the discriminant of T until we obtain an order which is p-maximal for every p, i.e. the maximal order. In practice however, even when the discriminant has square factors, $\mathbb{Z}[\theta]$ is quite often p-maximal for a number of primes p, and it is time consuming to have to compute $\mathcal{O}'$ as in Theorem 6.1.3 just to notice that $\mathcal{O}' = \mathbb{Z}[\theta]$, i.e. that $\mathbb{Z}[\theta]$ is p-maximal. Fortunately, there is a simple and important criterion due to Dedekind which allows us to decide, without the more complicated computations explained in the next section, whether $\mathbb{Z}[\theta]$ is p-maximal or not for prime numbers p, and if it is not, it will give us a larger order, which of course may still not be p-maximal.

It must be emphasized that this will work *only* for $\mathbb{Z}[\theta]$, or for any order $\mathcal{O}$ containing $\mathbb{Z}[\theta]$ with $[\mathcal{O} : \mathbb{Z}[\theta]]$ prime to p, but not for an order which has already been enlarged for the prime p itself.

This being said the basic theorem that we will prove, of which Dedekind's criterion is a special case, is as follows.

Theorem 6.1.4 (Dedekind). *Let $K = \mathbb{Q}(\theta)$ be a number field, $T \in \mathbb{Z}[X]$ the monic minimal polynomial of θ and let p be a prime number. Denote by $\bar{}$ reduction modulo p (in $\mathbb{Z}$, $\mathbb{Z}[X]$ or $\mathbb{Z}[\theta]$). Let*

$$\overline{T}(X) = \prod_{i=1}^{k} \overline{t}_i(X)^{e_i}$$

be the factorization of $T(X)$ *modulo* p *in* $\mathbb{F}_p[X]$, *and set*

$$g(X) = \prod_{i=1}^{k} t_i(X)$$

where the $t_i \in \mathbb{Z}[X]$ *are arbitrary monic lifts of* $\overline{t}_i$. *Then*

(1) *The* p*-radical* I_p *of* $\mathbb{Z}[\theta]$ *at* p *is given by*

$$I_p = p\mathbb{Z}[\theta] + g(\theta)\mathbb{Z}[\theta] \ .$$

In other words, $x = A(\theta) \in I_p$ *if and only if* $\overline{g} \mid \overline{A}$.

(2) *Let* $h(X) \in \mathbb{Z}[X]$ *be a monic lift of* $\overline{T}(X)/\overline{g}(X)$ *and set*

$$f(X) = (g(X)h(X) - T(X))/p \in \mathbb{Z}[X] \ .$$

Then $\mathbb{Z}[\theta]$ *is* p*-maximal if and only if*

$$(\overline{f}, \overline{g}, \overline{h}) = 1 \quad \text{in } \mathbb{F}_p[X] \ .$$

(3) *More generally, let* $\mathcal{O}'$ *be the order given by Theorem 6.1.3 when we start with* $\mathcal{O} = \mathbb{Z}[\theta]$. *Then, if* U *is a monic lift of* $\overline{T}/(\overline{f}, \overline{g}, \overline{h})$ *to* $\mathbb{Z}[X]$ *we have*

$$\mathcal{O}' = \mathbb{Z}[\theta] + \frac{1}{p}U(\theta)\mathbb{Z}[\theta]$$

and if $m = \deg(\overline{f}, \overline{g}, \overline{h})$, *then* $[\mathcal{O}' : \mathbb{Z}[\theta]] = p^m$, *hence* $\mathrm{disc}(\mathcal{O}') = \mathrm{disc}(T)/p^{2m}$.

Proof of (1). $p \in I_p$ trivially, and since the exponents e_i are at most equal to $n = [K : \mathbb{Q}] = \deg(T)$, we have $\overline{T} \mid \overline{g}^n$ hence $g^n(\theta) \equiv 0 \pmod{p\mathbb{Z}[\theta]}$ so $g(\theta) \in I_p$, thus proving that $I_p \supset p\mathbb{Z}[\theta] + g(\theta)\mathbb{Z}[\theta]$.

Now the minimal polynomial over $\mathbb{F}_p$ of θ in $\mathbb{Z}[\theta]/p\mathbb{Z}[\theta]$ (which is not a field in general) is clearly the polynomial $\overline{T}$. Indeed, it clearly divides $\overline{T}$, but it is of degree at least n since $1, \theta, \ldots, \theta^{n-1}$ are $\mathbb{F}_p$-linearly independent.

Conversely let $x \in I_p$. Then $x = A(\theta)$ for $A \in \mathbb{Z}[X]$, and so there exists an integer m such that $x^m \equiv 0 \pmod{p\mathbb{Z}[\theta]}$, in other words $\overline{A}^m(\theta) = 0$ in $\mathbb{Z}[\theta]/p\mathbb{Z}[\theta]$. Hence $\overline{T} \mid \overline{A}^m$. Since $e_i \geq 1$ for all i, this implies that $\overline{t}_i \mid \overline{A}^m$ hence $\overline{t}_i \mid \overline{A}$ since t_i is irreducible in $\mathbb{F}_p[X]$, and since the $\overline{t}_i$ are pairwise coprime, we get $\overline{g} \mid \overline{A}$ which means that $x \in p\mathbb{Z}[\theta] + g(\theta)\mathbb{Z}[\theta]$ thus proving (1).

Since $\overline{T}$ is the minimal polynomial of θ in $\mathbb{Z}[\theta]/p\mathbb{Z}[\theta]$, it is clear that (2) follows from (3).

Let us now prove (3). Recall that $\mathcal{O}' = \{x \in K | xI_p \subset I_p\}$. From (1) we have that $x \in \mathcal{O}'$ if and only if $xp \in I_p$ and $xg(\theta) \in I_p$. Since $I_p \subset \mathbb{Z}[\theta]$, $xp \in I_p$ implies that

$$x = A_1(\theta)/p$$

where $A_1 \in \mathbb{Z}[X]$. Part (3) of the theorem will immediately follow from the following lemma.

Lemma 6.1.5. *Let* $x = A_1(\theta)/p$ *with* $A_1 \in \mathbb{Z}[X]$. *Then*

(1) $xp \in I_p$ *if and only if*

$$\overline{g} \mid \overline{A_1} \ .$$

(2) *Let* $\overline{k} = \overline{g}/(\overline{f}, \overline{g})$, *where (here as elsewhere in this section)* k *is implicitly considered to be a monic lift of* $\overline{k}$ *to* $\mathbb{Z}[X]$. *Then* $xg(\theta) \in I_p$ *if and only if*

$$\overline{hk} \mid \overline{A_1} \ .$$

Proof of the Lemma. Part (1) of the lemma is an immediate consequence of part (1) of the theorem. Let us prove part (2).

From part (1) of the theorem, $xg(\theta) \in I_p$ if and only if there exist polynomials A_2 and A_3 in $\mathbb{Z}[X]$ such that

$$A_1(\theta)g(\theta) = p(pA_2(\theta) + g(\theta)A_3(\theta)) \ ,$$

and since T is the minimal polynomial of θ, this is true if and only if there exists $A_4 \in \mathbb{Z}[X]$ such that

$$A_1(X)g(X) = p^2A_2(X) + pg(X)A_3(X) + A_4(X)T(X) \ .$$

For the rest of this proof, we will work only with polynomials (in $\mathbb{Z}[X]$ or $\mathbb{F}_p[X]$), and not any more in K.

Reducing modulo p, the above equation implies that $\overline{A_1} = \overline{A_4}\overline{h}$. Hence write

$$A_1 = hA_4 + pA_5$$

with $A_5 \in \mathbb{Z}[X]$. We have that $xg(\theta) \in I_p$ if and only if there exist polynomials $A_i \in \mathbb{Z}[X]$ such that

$$(gh - T)A_4 = p^2A_2 + pg(A_3 - A_5) \ ,$$

hence if and only if there exist A_i such that

$$fA_4 = pA_2 + gA_6 \ .$$

This last condition is equivalent to $\overline{g} \mid \overline{fA_4}$ so to $\overline{k} \mid \overline{A_4}$ where $\overline{k} = \overline{g}/(\overline{f}, \overline{g})$, and this is equivalent to the existence of A_7 and A_8 in $\mathbb{Z}[X]$ such that $A_4 = kA_7 + pA_8$.

To sum up, we see that if $x = A_1(\theta)/p$, then $xg(\theta) \in I_p$ if and only if there exist polynomials A_5, A_7 and A_8 in $\mathbb{Z}[X]$ such that

$$A_1 = hkA_7 + p(hA_8 + A_5) \ ,$$

and this is true if and only if there exist $A_9 \in \mathbb{Z}[X]$ such that $A_1 = hkA_7 + pA_9$ or equivalently $\overline{hk} \mid \overline{A_1}$, thus proving the lemma. □

We can now prove part (3) of the theorem. From the lemma, we have that $x = A_1(\theta)/p \in \mathcal{O}'$ if and only if both $\overline{g}$ and $\overline{hk}$ divide $\overline{A_1}$ in the PID $\mathbb{F}_p[X]$, hence if and only if the lowest common multiple (lcm) of $\overline{g}$ and $\overline{hk}$ divides A_1. Since in any PID $\operatorname{lcm}(x,y) = xy/(x,y)$ and $\operatorname{lcm}(zx,zy) = z\operatorname{lcm}(x,y)$, we have

$$\operatorname{lcm}(\overline{g},\overline{hk}) = \overline{k}\operatorname{lcm}(\gcd(\overline{f},\overline{g}),\overline{h}) = \frac{\overline{g}}{(\overline{f},\overline{g})}\frac{\overline{h}(\overline{f},\overline{g})}{(\overline{f},\overline{g},\overline{h})} = \frac{\overline{T}}{(\overline{f},\overline{g},\overline{h})} = \overline{U}$$

thus proving that $\mathcal{O}' = \mathbb{Z}[\theta] + (U(\theta)/p)\mathbb{Z}[\theta]$. Now it is clear that a system of representatives of $\mathcal{O}'$ modulo $\mathbb{Z}[\theta]$ is given by $A(\theta)U(\theta)/p$ where A runs over uniquely chosen representatives in $\mathbb{Z}[X]$ of polynomials in $\mathbb{F}_p[X]$ such that $\deg(A) < \deg(T) - \deg(U) = m$, thus finishing the proof of the theorem. □

An important remark is that the proof of this theorem is *local* at p, in other words we can copy it essentially verbatim if we everywhere replace $\mathbb{Z}[\theta]$ by any overorder $\mathcal{O}$ of $\mathbb{Z}[\theta]$ such that $[\mathcal{O} : \mathbb{Z}[\theta]]$ is coprime to p. The final result is then that the new order enlarged at p is

$$\mathcal{O} + \frac{U(\theta)}{p}\mathcal{O} \ ,$$

and $[\mathcal{O}' : \mathcal{O}] = p^m$.

6.1.3 Outline of the Round 2 Algorithm

From the Pohst-Zassenhaus theorem it is easy to obtain an algorithm for computing the maximal order. We will of course use the Dedekind criterion to simplify the first steps for every prime p.

Let $K = \mathbb{Q}(\theta)$ be a number field, where θ is an algebraic integer. Let T be the minimal polynomial of θ. We can write $\operatorname{disc}(T) = df^2$, where d is either 1 or a fundamental discriminant. If $\mathbb{Z}_K$ is the maximal order which we are looking for, then the index $[\mathbb{Z}_K : \mathbb{Z}[\theta]]$ has only primes dividing f as prime divisors because of Proposition 4.4.4. We are going to compute $\mathbb{Z}_K$ by successive enlargements from $\mathcal{O} = \mathbb{Z}[\theta]$, one prime dividing f at a time. For every p dividing f we proceed as follows. By using Dedekind's criterion, we check whether $\mathcal{O}$ is p-maximal and if it is not we enlarge it once using Theorem 6.1.4 (3) applied to $\mathcal{O}$. If the new discriminant is not divisible by p^2, then we

are done, otherwise we compute $\mathcal{O}'$ as described in Theorem 6.1.3. If $\mathcal{O}' = \mathcal{O}$, then $\mathcal{O}$ is p-maximal and we are finished with the prime p, so we move on to the next prime, if any. (Here again we can start using Dedekind's criterion.) Otherwise, replace $\mathcal{O}$ by $\mathcal{O}'$, and use the method of Theorem 6.1.3 again. It is clear that this algorithm is valid and will lead quite rapidly to the maximal order. This algorithm was the second one invented by Zassenhaus for maximal order computations, and so it has become known as the round 2 algorithm (the latest and most efficient is round 4).

What remains is to explain how to carry out explicitly the different steps of the algorithm, when we apply Theorem 6.1.3.

First, θ is fixed, and all ideals and orders will be represented by their upper triangular HNF as explained in Section 4.7.2. We must explain how to compute the HNF of I_p and of $\mathcal{O}'$ in terms of the HNF of $\mathcal{O}$. It is simpler to compute in $R = \mathcal{O}/p\mathcal{O}$. To compute the radical of R, we note the following lemma:

Lemma 6.1.6. *If $n = [K : \mathbb{Q}]$ and if $j \geq 1$ is such that $p^j \geq n$, then the radical of R is equal to the kernel of the map $x \mapsto x^{p^j}$, which is the j^{th} power of the Frobenius homomorphism.*

Proof. It is clear that the map in question is the j^{th} power of the Frobenius homomorphism, hence talking about its kernel makes sense. By definition of the radical, it is clear that this kernel is contained in the radical. Conversely, let x be in the radical. Then x induces a nilpotent map defined by multiplication by x from R to R, and considering R as an $\mathbb{F}_p$-vector space, this means that the eigenvalues of this map in $\overline{\mathbb{F}_p}$ are all equal to 0. Hence, its characteristic polynomial must be X^n (since $n = \dim_{\mathbb{F}_p} R$), and by the Cayley-Hamilton theorem this shows that $x^n = 0$, and hence that $x^{p^j} = 0$, proving the lemma. □

Let $\omega_1, \ldots, \omega_n$ be the HNF basis of $\mathcal{O}$. Then it is clear that $\overline{\omega}_1, \ldots, \overline{\omega}_n$ is an $\mathbb{F}_p$-basis of R. For $k = 1, \ldots, n$, we compute $\overline{a}_{i,k}$ such that

$$\overline{\omega}_k^{p^j} = \sum_{i=1}^{n} \overline{a}_{i,k}\overline{\omega}_i \ ,$$

the left hand side being computed as a polynomial in θ by the standard representation algorithms, and the coefficients $\overline{a}_{i,k}$ being easily found inductively since an HNF matrix is triangular. Hence, if $\overline{A}$ is the matrix of the $\overline{a}_{i,k}$, the radical is simply the kernel of this matrix.

Hence, if we apply Algorithm 2.3.1, we will obtain a basis of $\overline{I}_p$, the radical of R, in terms of the standard representation. Since I_p is generated by pullbacks of a basis of $\overline{I}_p$ and $p\omega_1, \ldots, p\omega_n$, to obtain the HNF of I_p we apply the HNF reduction algorithm to the matrix whose columns are the standard representations of these elements.

Now that we have I_p, we must compute $\mathcal{O}'$. For this, we use the following lemma:

Lemma 6.1.7. *With the notations of Theorem 6.1.3, if U is the kernel of the map*

$$\alpha \longmapsto (\overline{\beta} \mapsto \overline{\alpha\beta})$$

from $\mathcal{O}$ to $\mathrm{End}(I_p/pI_p)$, then $O' = \frac{1}{p}U$.

Proof. Trivial and left to the reader. Note that $\mathrm{End}(I_p/pI_p)$ is considered as a $\mathbb{Z}$-module. □

Hence, we first need to find a basis of I_p/pI_p. There are two methods to do this. From the HNF reduction above, we know a basis of I_p, and it is clear that the image of this basis in I_p/pI_p is a basis of I_p/pI_p. The other method is as follows. We use only the $\mathbb{F}_p$-basis $\overline{\beta}_1, \ldots, \overline{\beta}_l$ of $\overline{I}_p$ found above. Using Algorithm 2.3.6, we can supplement this basis into a basis $\overline{\beta}_1, \ldots, \overline{\beta}_l, \overline{\beta}_{l+1}, \ldots, \overline{\beta}_n$ of $\mathcal{O}/p\mathcal{O}$, and then $\tilde{\beta}_1, \ldots, \tilde{\beta}_l, p\tilde{\beta}_{l+1}, \ldots, p\tilde{\beta}_n$ will be an $\mathbb{F}_p$-basis of I_p/pI_p, where $\tilde{}$ denotes reduction modulo pI_p, and β_i denotes any pull-back of $\overline{\beta}_i$ in $\mathcal{O}$. (Note that the basis which one obtains depends on the pull-backs used.)

This method for finding a basis of I_p/pI_p has the advantage of staying at the mod p level, hence avoids the time consuming Hermite reduction, so it is preferable.

Now that we have a basis of I_p/pI_p, the elementary matrices give us a basis of $\mathrm{End}(I_p/pI_p)$. Hence, we obtain explicitly the matrix of the map whose kernel is U, and it is a $n^2 \times n$ matrix. Algorithm 2.3.1 makes sense only over a field, so we must first compute the kernel $\overline{U}$ of the map from $\mathcal{O}/p\mathcal{O}$ into $\mathrm{End}(I_p/pI_p)$ which can be done using Algorithm 2.3.1. If $\overline{v}_1, \ldots, \overline{v}_k$ is the basis of this kernel, to obtain U, we apply Hermite reduction to the matrix whose column vectors are $v_1, \ldots, v_k, p\omega_1, \ldots, p\omega_n$. In fact, we can apply Hermite reduction modulo the prime p, i.e. take $D = p$ in Algorithm 2.4.8.

Finally, note that to obtain the $n^2 \times n$ matrix above, if the γ_i are a basis of I_p/pI_p one computes

$$\omega_k\overline{\gamma}_i = \sum_{1\le j\le n} a_{k,i,j}\overline{\gamma}_j \ ,$$

and k is the column number, while (i,j) is the row index. Unfortunately, in the round 2 algorithm, it seems unavoidable to use such large matrices. Note that to obtain the $a_{k,i,j}$, the work is much simpler if the matrix of the γ_j is triangular, and this is not the case in general if we complete the basis as explained above. On the other hand, this would be the case if we used the first method consisting of applying Hermite reduction to get the HNF of I_p itself. Tests must be made to see which method is preferable in practice.

6.1.4 Detailed Description of the Round 2 Algorithm

Using what we have explained, we can now give in complete detail the round 2 algorithm.

Algorithm 6.1.8 (Zassenhaus's Round 2). Let $K = \mathbb{Q}(\theta)$ be a number field given by an algebraic integer θ as root of its minimal monic polynomial T of degree n. This algorithm computes an integral basis $\omega_1 = 1, \omega_2, \ldots, \omega_n$ of the maximal order $\mathbb{Z}_K$ (as polynomials in θ) and the discriminant of the field. All the computations in K are implicitly assumed to be done using the standard representation of numbers as polynomials in θ.

1. [Factor discriminant of polynomial] Using Algorithm 3.3.7, compute $D \leftarrow \mathrm{disc}(T)$. Then using a factoring algorithm (see Chapters 8 to 10) factor D in the form $D = D_0F^2$ where D_0 is either equal to 1 or to a fundamental discriminant.
2. [Initialize] For $i = 1, \ldots, n$ set $\omega_i \leftarrow \theta^{i-1}$.
3. [Loop on factors of F] If $F = 1$, output the integral basis ω_i (which will be in HNF with respect to θ), compute the product G of the diagonal elements of the matrix of the ω_i, set $d \leftarrow D/G^2$, output the field discriminant d and terminate the algorithm. Otherwise, let p be the smallest prime factor of F.
4. [Factor modulo p] Using the mod p factoring algorithms of Section 3.4, factor T modulo p as $\overline{T} = \prod \overline{t}_i^{e_i}$ where the $\overline{t}_i$ are distinct irreducible polynomials in $\mathbb{F}_p[X]$ and $e_i > 0$ for all i. Set $\overline{g} \leftarrow \prod \overline{t}_i$, $\overline{h} \leftarrow \overline{T}/\overline{g}$, $f \leftarrow (gh - T)/p$, $\overline{Z} \leftarrow (\overline{f}, \overline{g}, \overline{h})$, $\overline{U} \leftarrow \overline{T}/\overline{Z}$ and $m \leftarrow \deg(Z)$.
5. [Apply Dedekind] If $m = 0$, then $\mathcal{O}$ is p-maximal so while $p \mid F$ set $F \leftarrow F/p$, then go to step 3. Otherwise, for $1 \le i \le m$, let v_i be the column vector of the components of $\omega_i U(\theta)$ on the standard basis $1, \theta, \ldots, \theta^{n-1}$ and set $v_{m+j} = p\omega_j$ for $1 \le j \le n$.

   Apply the Hermite reduction Algorithm 2.4.8 to the $n \times (n+m)$ matrix whose column vectors are the v_i. (Note that the determinant of the final matrix is known to divide D.) If H is the $n \times n$ HNF reduced matrix which we obtain, set for $1 \le i \le n$, $\omega_i \leftarrow H_i/p$ where H_i is the i-th column of H.
6. [Is the new order p-maximal?] If $p^{m+1} \nmid F$, then the new order is p-maximal so while $p \mid F$ set $F \leftarrow F/p$, then go to step 3.
7. [Compute radical] Set $q \leftarrow p$, and while $q < n$ set $q \leftarrow qp$. Then compute the $n \times n$ matrix $A = (a_{i,j})$ over $\mathbb{F}_p$ such that $\omega_j^q \equiv \sum_{1\le i\le n} a_{i,j}\omega_i$. Note that the matrix of the ω_i will stay triangular, so the $a_{i,j}$ are easy to compute.

   Finally, using Algorithm 2.3.1, compute a basis $\overline{\beta_1}, \ldots, \overline{\beta_l}$ of the kernel of the matrix A over $\mathbb{F}_p$ (this will be a basis of $I_p/p\mathcal{O}$).
8. [Compute new basis mod p] Using the known basis $\overline{\omega_1}, \ldots, \overline{\omega_n}$ of $\mathcal{O}/p\mathcal{O}$, supplement the linearly independent vectors $\overline{\beta_1}, \ldots, \overline{\beta_l}$ to a basis $\overline{\beta_1}, \ldots, \overline{\beta_n}$ of $\mathcal{O}/p\mathcal{O}$ using Algorithm 2.3.6.
9. [Compute big matrix] Set $\alpha_i \leftarrow \beta_i$ for $1 \le i \le l$, $\alpha_i \leftarrow p\beta_i$ for $l < i \le n$, where β_i is a lift to $\mathcal{O}$ of $\overline{\beta_i}$. Compute coefficients $c_{i,j,k} \in \mathbb{F}_p$ such that

$\omega_k \alpha_j \equiv \sum_{1 \le i \le n} c_{i,j,k} \alpha_i \pmod{p}$. Let C be the $n^2 \times n$ matrix over $\mathbb{F}_p$ such that $C_{(i,j),k} = c_{i,j,k}$.

10. [Compute new order] Using Algorithm 2.3.1, compute a basis $\gamma_1, \ldots \gamma_m$ for the kernel of C (these are vectors in $\mathbb{F}_p^n$, and m can be as large as n^2). For $1 \le i \le m$ let v_i be a lift of γ_i to $\mathbb{Z}^n$, and set $v_{m+j} = p\omega_j$ for $1 \le j \le n$. Apply the Hermite reduction Algorithm 2.4.8 to the $n \times (n+m)$ matrix whose column vectors are the v_i. (Note again that the determinant of the final matrix is known to divide D.) If H is the $n \times n$ HNF reduced matrix which we obtain, set for $1 \le i \le n$, $\omega_i' \leftarrow H_i/p$ where H_i is the i-th column of H.

11. [Finished with p?] If there exists an i such that $\omega_i' \neq \omega_i$, then for every i such that $1 \le i \le n$ set $\omega_i \leftarrow \omega_i'$ and go to step 7. Otherwise, $\mathcal{O}$ is p-maximal, so while $p \mid F$ set $F \leftarrow F/p$, and go to step 3.

This finishes our description of the round 2 algorithm. This algorithm seems complicated at first. Although it has been superseded by the round 4 algorithm, it is much simpler to implement and it performs very well. The major bottleneck is perhaps not where the reader expects it to be, i.e. in the handling of large matrices. It is, in fact, in the very first step which consists in factoring $\mathrm{disc}(T)$ in the form $D_0 F^2$. Indeed, as we will see in Chapter 10, factoring an 80 digit number takes a considerable amount of time, and factoring a 50 digit one is already not that easy. One can refine the methods given above to the case where one does not suppose p to be necessarily prime (see [Buc-Len]), but unfortunately this does *not* avoid finding the largest square dividing $\mathrm{disc}(T)$, which is apparently almost as difficult as factoring it completely.

6.2 Decomposition of Prime Numbers II

As we shall see, the general problem of decomposing prime numbers in an algebraic number field is closely related to the problem of computing the maximal order. Consequently, we have already given most of the theory and auxiliary algorithms that we will need. As we have already seen, the problem is as follows. Given a prime p and a p-maximal order $\mathcal{O}$, for example the maximal order $\mathbb{Z}_K$ itself, determine the maximal ideals $\mathfrak{p}_i$ and the exponents e_i such that

$$p\mathcal{O} = \prod_{i=1}^{g} \mathfrak{p}_i^{e_i} .$$

As usual $\mathcal{O}$ will be given by its HNF on a power basis $1, \theta, \ldots, \theta^{n-1}$, and we want the HNF basis of the $\mathfrak{p}_i$. The determinant of the corresponding matrix is equal to $\mathcal{N}(\mathfrak{p}_i) = p^{f_i}$ in the traditional notation. For practical applications, it will also be useful to have a two-element representation of the ideals $\mathfrak{p}_i$ (see Proposition 4.7.7).

In Theorem 4.8.13 we saw how to obtain this decomposition when p does not divide the index $[\mathcal{O} : \mathbb{Z}[\theta]]$. Hence we will concentrate on the case where p divides the index.

6.2.1 Newton Polygons

Historically the first method to deal with this problem is the so-called *Newton polygon method*. When it applies, it is very easy to use, but it must be stressed that it is not a general method. We will give a completely general method in the next section.

I am grateful to F. Diaz y Diaz and M. Olivier for the presentation of Newton polygons given here, which follows [Ore] and [Mon-Nar]. Essentially no proofs are given.

We may assume without loss of generality that the minimal polynomial $T(X)$ of θ is in $\mathbb{Z}[X]$ and is monic.

The first result tells us what survives of Theorem 4.8.13 in the case where p divides the index.

Proposition 6.2.1. *Let*

$$T(X) \equiv \prod_{i=1}^{g} \overline{T_i(X)}^{e_i} \pmod{p}$$

be the decomposition of T into irreducible factors in $\mathbb{F}_p[X]$, where the T_i are taken to be arbitrary monic lifts of $\overline{T_i(X)}$ in $\mathbb{Z}[X]$. Then

$$p\mathbb{Z}_K = \prod_{i=1}^{g} \mathfrak{a}_i ,$$

where

$$\mathfrak{a}_i = (p, T_i^{e_i}(\theta)) = p\mathbb{Z}_K + T_i^{e_i}(\theta)\mathbb{Z}_K$$

and the $\mathfrak{a}_i$ are pairwise coprime (i.e. $\mathfrak{a}_i + \mathfrak{a}_j = \mathbb{Z}_K$ for $i \neq j$). Furthermore, if n_i is the degree of T_i we have $\mathcal{N}(\mathfrak{a}_i) = p^{e_i n_i}$, and all prime ideals dividing $\mathfrak{a}_i$ are of residual degree divisible by n_i.

Proof. The proof follows essentially the same lines as that of Theorem 4.8.13. It is useful to also prove that the inverse of $\mathfrak{a}_i$ is given explicitly as

$$\mathfrak{a}_i^{-1} = (1, \prod_{j \neq i} T_j^{e_j}(\theta)/p)$$

(see Exercise 5). □

The problem is that the ideals $\mathfrak{a}_i$ are not necessarily of the form $\mathfrak{p}_i^{e_i}$ as in Theorem 4.8.13 (the reader can also check via examples that it would not