

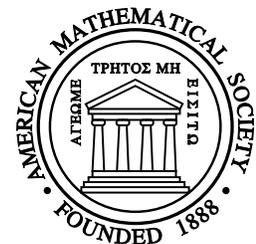
ENTANGLED RADICALS

Hendrik W. Lenstra jr.
 Universiteit Leiden

Colloquium Lectures



American Mathematical Society
 112th Annual Meeting
 San Antonio, January 12–15, 2006



Profinite Fibonacci numbers

The picture on the title page, which was created by W. J. Palenstijn, is taken from the author's paper *Profinite Fibonacci numbers*, Nieuw Archief voor Wiskunde (5) **6** (2005), 297–300. It illustrates several features of the ring $\hat{\mathbf{Z}}$ of profinite integers, as defined in Section 7 of the present lecture notes. The *Fibonacci function*, written $z \mapsto F_z$, is the unique continuous function $\hat{\mathbf{Z}} \rightarrow \hat{\mathbf{Z}}$ satisfying $F_0 = 0$, $F_1 = 1$, $F_{z+2} = F_{z+1} + F_z$ for all $z \in \hat{\mathbf{Z}}$; thus F_n is, for each positive integer n , equal to the familiar n th Fibonacci number. Each profinite integer z has a unique representation as an infinite sum $\sum_{i>0} c_i i!$, with each c_i belonging to $\{1, 2, \dots, i\}$, and we represent z by the number $\sum_{i>0} c_i / (i + 1)!$ in the unit interval. The graph $\{(z, F_z) : z \in \hat{\mathbf{Z}}\}$ is correspondingly represented as a subset of the unit square. Successive approximations to the graph are shown in orange, red, and brown. Intersecting the diagonal, shown in green, with the graph, one finds the *fixed points* of the Fibonacci function, of which there are eleven: the fixed point 0, and, for each $a \in \{1, 5\}$ and $b \in \{-5, -1, 0, 1, 5\}$, a unique fixed point $z_{a,b}$ that satisfies $z_{a,b} \equiv a \pmod{6^k}$, $z_{a,b} \equiv b \pmod{5^k}$ for all positive integers k ; one has $z_{1,1} = 1$ and $z_{5,5} = 5$. There are two clusters of three fixed points each that are indistinguishable in the precision used. One of these clusters is resolved by a sequence of three blow-ups, with a total magnification factor of $14 \times 16 \times 18 = 4032$. The graph of the function $z \mapsto -z$, shown in blue, enables the viewer to check the formula $F_{-z} = (-1)^{z-1} F_z$. The yellow squares contain the curve $\{(z, w) \in \hat{\mathbf{Z}} \times \hat{\mathbf{Z}} : z \cdot (w + 1) = 1\}$. Its projection on the horizontal axis equals the group $\hat{\mathbf{Z}}^*$ of units of $\hat{\mathbf{Z}}$, which plays an important role in Section 13 of the lecture notes.

Abstract

These lectures are concerned with the algebra of radicals, that is, with algebraic expressions involving root signs. The first two lectures address the structure of the field that one obtains by adjoining, to a given base field of characteristic zero, all elements in an algebraic closure that have a positive power lying in the base field. The ulterior motive is creating the possibility of correctly computing with radical expressions in computer algebra systems. The third lecture is correspondingly devoted to algorithmic issues. It will become clear that much theoretical and algorithmic work remains to be done.

Entangled radicals

Hendrik W. Lenstra jr.

Mathematisch Instituut, Universiteit Leiden,

Postbus 9512, 2300 RA Leiden, The Netherlands

1. Introduction	4–5
2. Entanglement	5–6
3. Artin’s conjecture	6–8
4. The group of radicals	8–10
5. A modified group ring	10–12
6. Radicals over the field of rational numbers	12–14
7. Profinite integers	14–15
8. Profinite groups	15–16
9. Infinite Galois theory	16–17
10. Galois algebras	17–19
11. The ring of radicals	19–20
12. The algebra of the ring of radicals	21–22
13. The radical Galois group	22–24
14. The decomposition algebra	24–25
15. Small subgroups of the group of radicals	25–26
16. Finite étale algebras	27
17. Fields generated by finitely many radicals	28–29
18. Algorithmic issues	29–30
19. Computing with radicals	31–32
20. Polynomial-time algorithms	32–33
21. Nested radicals	33–34
22. Algorithms for large number fields	35–36
23. Abelian and solvable fields	36–37
Acknowledgments	38
References	38–39

1. Introduction

Let K be a field of characteristic zero, and denote by Ω an algebraically closed field that contains K . For example, K and Ω may be the fields \mathbf{Q} and \mathbf{C} of rational and complex numbers, respectively. Write K^* for the multiplicative group of non-zero elements of K , and $\sqrt{K^*}$ for the group of *radicals* over K , i.e., the subgroup $\{a \in \Omega^* : a^n \in K^* \text{ for some positive integer } n\}$ of Ω^* . We denote by $K(\sqrt{K^*})$ the subfield of Ω obtained by adjoining $\sqrt{K^*}$ to K .

The structure of $K(\sqrt{K^*})$, as an extension field of K , is independent of the choice of Ω , and the question poses itself to “understand” this structure in terms of the field K itself. It is to this question that the present series of lectures is devoted.

There are several ways in which the question can be made precise. One may interpret it in a practical way, from the point of view of the designer of a computer algebra system who wishes to enable his customers to do computations with radicals. Current computer algebra systems allow arithmetic in several types of fields, such as number fields and function fields of varieties, as well as various sorts of completions of these fields. Over all these fields one may wish to be able to compute with radicals.

Interpreting the question in a more theoretical manner, one may investigate the field extension $K(\sqrt{K^*})$ of K from the point of view of abstract algebra. It is a *Galois extension*, with a Galois group $\text{Gal}(K(\sqrt{K^*})/K)$ that, by definition, consists of all field automorphisms of $K(\sqrt{K^*})$ that are the identity on K . Can one identify this Galois group? Each $\sigma \in \text{Gal}(K(\sqrt{K^*})/K)$ restricts to a group automorphism of the multiplicative group $\sqrt{K^*}$ that is the identity on K^* ; denoting the group of such group automorphisms by $\text{Aut}_{K^*}\sqrt{K^*}$, we see that $\text{Gal}(K(\sqrt{K^*})/K)$ may be viewed as a subgroup of $\text{Aut}_{K^*}\sqrt{K^*}$. Which subgroup is it?

Our attitude will be largely theoretical, but with constant attention for practical issues. It turns out that the two interpretations of our question are not as unrelated as they may seem. Determining $\text{Gal}(K(\sqrt{K^*})/K)$ is in fact tantamount to understanding the structure of $K(\sqrt{K^*})$ as a field extension of K . Acquiring such understanding is the main mathematical obstacle that the designer of a system for computing in $K(\sqrt{K^*})$ needs to overcome. We do not claim that there are no other mathematical difficulties to be faced, and the non-mathematical ones—involving, among others, the psychology and the expectations of the user—may in fact be the hardest of all. Thus, while we are able to offer a satisfactory solution to our theoretical problem, we are quite far from offering ready-made solutions for

the practical one. We do report on algorithmic work that has been done, and formulate a number of problems that merit further investigation.

2. Entanglement

Let the notation be as in the introduction. Since the field $K(\sqrt{K^*})$ is algebraic over K , it is equal to the *ring* generated by $\sqrt{K^*}$ over K . It follows that each element of $K(\sqrt{K^*})$ can be written as a sum of finitely many elements of $\sqrt{K^*}$. Both from a theoretical and from a practical point of view, this is a natural representation to use. It leads to two fundamental problems.

The first problem is of a purely multiplicative nature: how does one represent the elements of $\sqrt{K^*}$ itself in such a manner that the group operations in $\sqrt{K^*}$ are readily performed? Using the symbol \sqrt{a} to represent a zero of the polynomial $X^2 - a$ is inherently ambiguous, and cavalier use of the “rule” $\sqrt{a} \cdot \sqrt{b} = \sqrt{ab}$ leads to contradictions such as

$$1 = \sqrt{1} \cdot \sqrt{1} = \sqrt{1 \cdot 1} = \sqrt{-1 \cdot -1} = \sqrt{-1} \cdot \sqrt{-1} = -1.$$

For “real” radicals there are familiar positivity conventions, but they are not easily extendable to general K .

The second problem is of an additive nature: how can we tell whether two finite sums of elements of $\sqrt{K^*}$ represent the same element of $K(\sqrt{K^*})$? With a suitable interpretation of the square-roots, one has $\sqrt{2} + \sqrt{2} = \sqrt{8}$ (and with another one, one has $\sqrt{2} + \sqrt{2} = 0$), and generally, if $\alpha, \beta \in \sqrt{K^*}$ belong to the same coset modulo the subgroup K^* , then one has $\alpha + \beta \in \sqrt{K^*}$ or $\alpha + \beta = 0$. Thus, it is natural to restrict to sums $\sum_i \alpha_i$ of finitely many elements $\alpha_i \in \sqrt{K^*}$ with the restriction that no two α_i belong to the same coset modulo K^* . It is, however, naïve to expect that representations by such sums are unique. Non-uniqueness has two sources: first, trivial linear relations satisfied by roots of unity, and, second, the truly *entangled radicals*.

We give examples of both types with $K = \mathbf{Q}$. If $\zeta \in \sqrt{\mathbf{Q}^*}$ has order 5, then one has

$$1 + \zeta + \zeta^2 + \zeta^3 + \zeta^4 = \frac{1 - \zeta^5}{1 - \zeta} = 0,$$

while $1, \zeta, \zeta^2, \zeta^3, \zeta^4$ belong to different cosets of $\sqrt{\mathbf{Q}^*}$ modulo \mathbf{Q}^* . Additive relations of this nature are relatively harmless. Over other fields, roots of unity may satisfy more subtle linear relations.

Examples of truly entangled radicals over \mathbf{Q} are given by the “identities”

$$1 + \sqrt{-1} = \sqrt[4]{-4}, \quad \zeta - \zeta^2 - \zeta^3 + \zeta^4 = \sqrt{5},$$

with ζ as above; more precisely, $1 + \sqrt{-1}$ is, for any zero $\sqrt{-1}$ of $X^2 + 1$, a zero of $X^4 + 4$ (although not every zero of $X^4 + 4$ is of that form), and as ζ ranges over all four elements of $\sqrt{\mathbf{Q}^*}$ of order 5 (i.e., over all zeroes of $X^4 + X^3 + X^2 + X + 1$), then $\zeta - \zeta^2 - \zeta^3 + \zeta^4$ ranges over all zeroes of $X^2 - 5$. Again, $1, \sqrt{-1}, \sqrt[4]{-4}$ belong to different cosets modulo \mathbf{Q}^* , and so do $\zeta, \zeta^2, \zeta^3, \zeta^4, \sqrt{5}$.

Coming to terms with relations of the latter type is the main challenge of the subject.

The consequences of entangled radicals in number theory can be surprising. Here are two examples: *the only positive integer n for which $n^4 + 4^n$ is prime is $n = 1$* , and: *for every prime number p with $p \equiv 1 \pmod{4}$, the number $(p^p - 1)/(p - 1)$ is composite*. Results of this nature are often named after the 19th century French mathematician Antoine Aurifeuille.

3. Artin’s conjecture

A characteristic manifestation of entangled radicals is encountered in *Artin’s conjecture* on primes with prescribed primitive roots, which we discuss by way of intermezzo.

Let a be an integer that is not a proper power. What is the probability for a prime number to admit a as a primitive root? More precisely, does $\#\{p : p \text{ prime, } p \leq x, \text{ and } p \text{ has } a \text{ as a primitive root}\}$, divided by the total number of primes that are at most x , tend to a limit as $x \rightarrow \infty$, and what is the value of this limit? Emil Artin, in 1927, conjectured that the limit is, independently of a , equal to the number

$$A = \prod_{l \text{ prime}} \left(1 - \frac{1}{l(l-1)}\right) \doteq 0.373955814$$

(see [33]), which is now called *Artin’s constant*.

The reasoning that led Artin to his conjecture is as follows. Let p be a prime number not dividing a . By definition, a is a primitive root modulo p if and only if the group \mathbf{F}_p^* of non-zero residues modulo p is generated by $(a \pmod{p})$; this is clearly equivalent to the index $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle)$ of the subgroup generated by $(a \pmod{p})$ being non-zero modulo l for each prime number l .

Now fix, in addition to a , a prime number l . What is the probability for a prime number p to satisfy $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \equiv 0 \pmod{l}$? First, one needs to have $\#\mathbf{F}_p^* \equiv 0 \pmod{l}$

or, equivalently, $p \equiv 1 \pmod{l}$. Since the prime numbers different from l are equidistributed over the $l - 1$ non-zero residue classes modulo l , the probability for this to happen equals $1/(l-1)$. Next, for $p \equiv 1 \pmod{l}$, one has $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \equiv 0 \pmod{l}$ if and only if $(a \pmod{p})$ belongs to the subgroup \mathbf{F}_p^{*l} of l th powers of the cyclic group \mathbf{F}_p^* ; since that subgroup has index l , it seems plausible that it contains $(a \pmod{p})$ with probability $1/l$. Thus, the overall probability for a prime p to satisfy $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \equiv 0 \pmod{l}$ appears to be $1/(l(l-1))$. This can in fact be rigorously proved, by means of Chebotarëv's density theorem [22] and the fact that the splitting field of $X^l - a$ over \mathbf{Q} has degree equal to $l(l-1)$.

It follows that, for a fixed prime number l , the probability for a prime number p to satisfy $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \not\equiv 0 \pmod{l}$ equals $1 - 1/(l(l-1))$. It seems reasonable to assume that, for different l , the “events” $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \not\equiv 0 \pmod{l}$ are “independent”, so that the probability for p to satisfy this condition for *all* prime numbers l would appear to be given by the infinite product that defines Artin's constant. The latter part of the argument, in particular the passage to *infinitely* many l , is not easy to justify rigorously, and this is why Artin pronounced merely a *conjecture*.

Computations by D. H. Lehmer and E. Lehmer in 1957 exhibited, for many values of a , a discrepancy with Artin's prediction [19]. When Artin learnt of this, he realized his mistake: he had overlooked that, due to entanglement of radicals, the conditions $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \not\equiv 0 \pmod{l}$ may *not* be independent. Consider for example $a = 5$. If a prime number $p \neq 5$ satisfies $(\mathbf{F}_p^* : \langle 5 \pmod{p} \rangle) \equiv 0 \pmod{5}$, then \mathbf{F}_p^* has an element ζ of order 5; the identity $\zeta - \zeta^2 - \zeta^3 + \zeta^4 = \sqrt{5}$ that we saw above, remains valid under “reduction modulo p ” and gives rise to the equality $(\zeta - \zeta^2 - \zeta^3 + \zeta^4)^2 = (5 \pmod{p})$, so that 5 is a square in \mathbf{F}_p and consequently $(\mathbf{F}_p^* : \langle 5 \pmod{p} \rangle)$ is *even*. In other words, one has

$$(\mathbf{F}_p^* : \langle 5 \pmod{p} \rangle) \not\equiv 0 \pmod{2} \Rightarrow (\mathbf{F}_p^* : \langle 5 \pmod{p} \rangle) \not\equiv 0 \pmod{5}.$$

It follows that the condition for $l = 5$ should not be accounted for. The conjectural probability for a prime p to have 5 as a primitive root now becomes

$$\prod_{l \text{ prime}, l \neq 5} \left(1 - \frac{1}{l(l-1)}\right) = \frac{20}{19} \cdot A \doteq 0.393637699.$$

This is in much better agreement with Lehmer's calculations.

For general a , accounting for the dependencies between the conditions $(\mathbf{F}_p^* : \langle a \pmod{p} \rangle) \not\equiv 0 \pmod{l}$ leads to the formulation of a corrected conjecture, in which Artin's constant

A is replaced by $E_a \cdot A$ for a certain rational error factor E_a depending on a . The corrected form of the conjecture was shown by Hooley [15] to be a consequence of the generalized Riemann hypothesis. The original form of the conjecture is known to be false for infinitely many values of a , such as $a = 21$.

There is a fair amount of literature on the correction factors required in Artin's conjecture and its many generalizations. Invariably, these factors are given by formulae that are much simpler than one seems to have a right to expect. An explanation is furnished by the theory that we are about to describe; this is one of the subjects that the special session accompanying the present series of lectures is devoted to (cf. [29]).

4. The group of radicals

Let K be a field of characteristic zero. In Section 2, we identified two problems in understanding the field extension $K(\sqrt{K^*})$ of K , a multiplicative one and an additive one. In the present section we address the multiplicative problem.

For a multiplicatively written abelian group A and an integer n , we write $A[n]$ and A^n , respectively, for the kernel and the image of the group homomorphism $A \rightarrow A$ sending each $a \in A$ to a^n . We say that A is *torsion* if $A = \bigcup_{n \geq 1} A[n]$.

The group $\sqrt{K^*}$ has clearly the following properties: (i) it is an abelian group containing K^* as a subgroup, and the group $\sqrt{K^*}/K^*$ is torsion; (ii) for each positive integer n , one has $\#\sqrt{K^*}[n] = n$ and $K^* \subset \sqrt{K^*}^n$. Our first observation is that $\sqrt{K^*}$ is to a certain extent uniquely determined by these properties. In fact, one may start from any abelian group B that shares with K^* the property that every finite subgroup is cyclic, and construct a corresponding group \sqrt{B} in a purely multiplicative way.

Theorem. *Let B be an abelian group such that every finite subgroup of B is cyclic. Then there exists an abelian group A with the following properties:*

- (i) *A contains B as a subgroup, and the group A/B is torsion;*
- (ii) *for each positive integer n , one has $\#A[n] = n$ and $B \subset A^n$.*

Moreover, A is uniquely determined up to isomorphism, in the following sense: if $\phi: B \rightarrow B'$ is an isomorphism of B with a group B' , and A' is an abelian group that has the same properties with B' in the role of B , then ϕ can be extended to a group isomorphism $\psi: A \rightarrow A'$.

The proof of this result is an exercise in abelian group theory (cf. [24]). We shall write \sqrt{B} for a group A of “radicals” of B of which the existence is asserted in the theorem.

If E denotes a *divisible hull* of B (see [11, Section 24]), then there is an isomorphism $\sqrt{B} \cong E \oplus \bigoplus_p (\mathbf{Z}[1/p]/\mathbf{Z})$ sending each $b \in B$ to $(b, 0)$; here p ranges over the set of prime numbers for which $\#B[p] = 1$, and $\mathbf{Z}[1/p]$ denotes the additive group of rational numbers of which the denominator is a power of p .

It is to be emphasized that \sqrt{B} is *not* “uniquely unique”; that is, the isomorphism ψ of which the existence is asserted in the theorem, is *not* uniquely determined by ϕ , except in the case $B = A$. Put in another way, there is no reasonable way of making $\sqrt{}$ into a *functor*. This is the reason why we do not call \sqrt{B} *the* group of radicals of B . The reader may note the similarity with the “non-unique” uniqueness of algebraic closures of fields.

For a field K of characteristic zero, the notation $\sqrt{K^*}$ has now several meanings: the one given in the introduction, and the one just defined, with $B = K^*$. In the sequel, we shall adhere to the latter meaning: given K^* , we choose a group $\sqrt{K^*}$ of radicals of K^* once and for all, and we think of it purely as a multiplicatively written abelian group containing K^* . By $\text{Aut}_{K^*}\sqrt{K^*}$ we denote the group of those group automorphisms of $\sqrt{K^*}$ that are the identity on K^* . For an algebraically closed field Ω containing K , we are interested in *embeddings* $\sqrt{K^*} \rightarrow \Omega^*$, that is, in group homomorphisms that are the identity on K^* . Such embeddings exist, and they all have the same image, namely the group $\sqrt{K^*}$ as defined in the introduction. It follows that for any two embeddings $s, t: \sqrt{K^*} \rightarrow \Omega^*$, there is a unique element $\sigma \in \text{Aut}_{K^*}\sqrt{K^*}$ such that $s = t \circ \sigma$. The subfield generated by the image of $\sqrt{K^*}$ in Ω is independent of the choice of the embedding, and we just write $K(\sqrt{K^*})$ for this field. One should note, however, that the restriction map $\text{Gal}(K(\sqrt{K^*})/K) \rightarrow \text{Aut}_{K^*}\sqrt{K^*}$ does depend on the embedding; if we write ρ_s for the restriction map induced by an embedding s , then for all $\tau \in \text{Gal}(K(\sqrt{K^*})/K)$ one has $s \circ \rho_s(\tau) = \tau \circ s$, and the restriction map ρ_t induced by an embedding $t = s \circ \sigma^{-1}$ (with $\sigma \in \text{Aut}_{K^*}\sqrt{K^*}$) is then given by $\rho_t(\tau) = \sigma \cdot \rho_s(\tau) \cdot \sigma^{-1}$. Thus, $\text{Gal}(K(\sqrt{K^*})/K)$ is, when viewed as a subgroup of $\text{Aut}_{K^*}\sqrt{K^*}$, *a priori* only well-defined up to conjugation.

As an example we consider the case $K = \mathbf{Q}$. By unique prime factorization, there is a group isomorphism

$$(\mathbf{Z}/2\mathbf{Z}) \oplus \bigoplus_{p \text{ prime}} \mathbf{Z} \rightarrow \mathbf{Q}^*$$

sending $(m \bmod 2, (m(p))_{p \text{ prime}})$ to $(-1)^m \cdot \prod_{p \text{ prime}} p^{m(p)}$. A group of radicals of the first group is given by $(\mathbf{Q}/2\mathbf{Z}) \oplus \bigoplus_{p \text{ prime}} \mathbf{Q}$. For $\sqrt{\mathbf{Q}^*}$ we can now take a multiplicatively written copy of the latter group; we shall write $(-1)^m \cdot \prod_{p \text{ prime}} p^{m(p)}$ for the element of $\sqrt{\mathbf{Q}^*}$ corresponding to an element $(m \bmod 2\mathbf{Z}, (m(p))_{p \text{ prime}})$ of $(\mathbf{Q}/2\mathbf{Z}) \oplus \bigoplus_{p \text{ prime}} \mathbf{Q}$.

This notation for elements of $\sqrt{\mathbf{Q}^*}$ avoids use of the radical sign $\sqrt{}$, which represents a psychological advantage. Of course, the difficulties that use of the $\sqrt{}$ -sign gives rise to, will resurface if one makes the mistake of thinking that there is a way of defining a map $\mathbf{Q} \times \sqrt{\mathbf{Q}^*} \rightarrow \sqrt{\mathbf{Q}^*}$, $(q, r) \mapsto r^q$, that satisfies the module axioms (namely: $(r_1 r_2)^q = r_1^q r_2^q$, $r^{q_1 + q_2} = r^{q_1} r^{q_2}$, $r^{q_1 q_2} = (r^{q_1})^{q_2}$, $r^1 = r$).

It may for general K not be as easy as for $K = \mathbf{Q}$ to describe the group $\sqrt{K^*}$ and to represent its elements in an explicit manner. In computational circumstances, one restricts to subgroups A of $\sqrt{K^*}$ that contain K^* as a subgroup of finite index, and the elements of such groups allow easy representations (cf. Section 15 below). In any case, the theorem shows that “understanding” the group $\sqrt{K^*}$ is a purely multiplicative issue. We shall simply assume that $\sqrt{K^*}$ is, as an abelian group containing K^* , somehow “under control”, and we address the problem of constructing the field extension $K(\sqrt{K^*})$ out of it.

5. A modified group ring

For any ring R and any multiplicatively written group A , the *group ring* $R[A]$ consists, by definition, of *finite* formal sums $\sum_{a \in A} r_a \cdot a$ with coefficients $r_a \in R$; finiteness of such a sum means that the set $\{a \in A : r_a \neq 0\}$ is finite. The group ring is a ring with addition $(\sum_{a \in A} r_a \cdot a) + (\sum_{a \in A} s_a \cdot a) = \sum_{a \in A} (r_a + s_a) \cdot a$ and multiplication $(\sum_{a \in A} r_a \cdot a) \cdot (\sum_{a \in A} s_a \cdot a) = \sum_{a \in A} (\sum_{b, c \in A, bc=a} r_b s_c) \cdot a$. It contains R as a subring, and its group of units contains A as a subgroup.

Let K be a field of characteristic zero, and let $\sqrt{K^*}$ be a group of radicals of K^* as defined in the previous section. As a first step towards constructing the field $K(\sqrt{K^*})$ out of the group $\sqrt{K^*}$, we form the group ring $K[\sqrt{K^*}]$. Note that the elements of $\sqrt{K^*}$ form a basis of this group ring when viewed as a vector space over K .

Let Ω be an algebraically closed field containing K . Clearly, any embedding $\sqrt{K^*} \rightarrow \Omega^*$ extends to a unique ring homomorphism $K[\sqrt{K^*}] \rightarrow \Omega$ that is the identity on K . Its image is the subring of Ω generated by K and the image of $\sqrt{K^*}$, and that subring is, as we observed in Section 2, all of the field $K(\sqrt{K^*})$. So, $K(\sqrt{K^*})$ may be identified with the group ring $K[\sqrt{K^*}]$ modulo some ideal \mathfrak{m} , namely, the kernel of the ring homomorphism; it is a *maximal* ideal, since $K[\sqrt{K^*}]/\mathfrak{m}$ is a field. We wish to control the set of maximal ideals of $K[\sqrt{K^*}]$ that can arise in this fashion, without reference to Ω . Evidently, independently of any choices, any of these maximal ideals contains all elements of the form $a \cdot 1 - 1 \cdot a$, with $a \in K^*$. Hence, we may as well factor out those elements, and replace $K[\sqrt{K^*}]$ by

the ring

$$K\{\sqrt{K^*}\} = K[\sqrt{K^*}]/\mathfrak{a},$$

where \mathfrak{a} denotes the ideal of $K[\sqrt{K^*}]$ generated by $\{a \cdot 1 - 1 \cdot a : a \in K^*\}$. This ring still contains K as a subring and $\sqrt{K^*}$ as a subgroup of its group of units. Also, with Ω as above, it is still true that any embedding $\sqrt{K^*} \rightarrow \Omega^*$ extends to a unique ring homomorphism $K[\sqrt{K^*}] \rightarrow \Omega$ that is the identity on K .

Proposition. *Any system of coset representatives for $\sqrt{K^*}$ modulo K^* forms a basis of $K\{\sqrt{K^*}\}$ when viewed as a vector space over K .*

The reader acquainted with tensor products can easily deduce this result from the equivalent description

$$K\{\sqrt{K^*}\} = K[\sqrt{K^*}] \otimes_{K[K^*]} K,$$

where $K[K^*]$ is again a group ring and the ring homomorphism $K[K^*] \rightarrow K$ sends a formal sum $\sum_{a \in K^*} r_a \cdot a$ to $\sum_{a \in K^*} r_a a$, the products $r_a a$ in the latter sum, as well as the sum itself, being evaluated in K .

The proposition gives us a way of uniquely representing the elements of $K\{\sqrt{K^*}\}$. It is equivalent to using finite sums $\sum_i \alpha_i$ of elements α_i of $\sqrt{K^*}$ belonging to pairwise different cosets modulo K^* .

We may now wonder whether the ring $K\{\sqrt{K^*}\}$ is actually a *field*. If this is the case, then $\{0\}$ is its only maximal ideal, so that we have $K\{\sqrt{K^*}\} \cong K(\sqrt{K^*})$. That would constitute a perfectly satisfactory description of the field $K(\sqrt{K^*})$, with no additive entanglement of radicals whatsoever occurring.

A classical chapter of field theory, referred to as *Kummer theory* (see [18, Section 6.8]), implies that $K\{\sqrt{K^*}\}$ is indeed a field isomorphic to $K(\sqrt{K^*})$ if K contains “all roots of unity”, in the sense that for all positive integers n one has $\#K^*[n] = n$. If that condition is satisfied, then $K\{\sqrt{K^*}\}$ is a Galois extension of K in the sense of (infinite) Galois theory for fields, with a Galois group that is abelian, and with canonical group isomorphisms

$$\text{Gal}(K\{\sqrt{K^*}\}/K) \cong \text{Aut}_{K^*} \sqrt{K^*} \cong \text{Hom}(K^*, T(\mathbf{G}_m)),$$

for a certain abelian group $T(\mathbf{G}_m)$ to be considered in Section 8.

There is actually a weaker condition under which $K\{\sqrt{K^*}\}$ is a field. It results from “Kneser theory” (see [16; 1]), which is a little less classical than Kummer theory. To state the result, it is convenient to fix an embedding of $\sqrt{K^*}$ in Ω^* , with Ω as before.

Theorem. *Let K be a field of characteristic zero, and let the ring $K\{\sqrt{K^*}\}$ be as defined above. Then the following statements are equivalent:*

- (i) $K\{\sqrt{K^*}\}$ is a field;
- (ii) $K\{\sqrt{K^*}\}$ is a domain;
- (iii) the ring homomorphism $K\{\sqrt{K^*}\} \rightarrow K(\sqrt{K^*})$ is an isomorphism;
- (iv) the group homomorphism $\text{Gal}(K(\sqrt{K^*})/K) \rightarrow \text{Aut}_{K^*}\sqrt{K^*}$ is an isomorphism;
- (v) the group K^* has an element of order 4 and, for each odd prime number p , an element of order p .

Condition (v) is the one occurring in Kneser theory. The equivalence between (iii) and (iv) illustrates that determining $\text{Gal}(K(\sqrt{K^*})/K)$ is tantamount to understanding the structure of $K(\sqrt{K^*})$ as a field extension of K , as we stated in the introduction. One difference between Kneser theory and Kummer theory is that the Galois group in (iv) is not necessarily abelian.

6. Radicals over the field of rational numbers

In this section we consider the case $K = \mathbf{Q}$. We begin by describing the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$.

As we saw in Section 4, the group $\sqrt{\mathbf{Q}^*}$ is the direct sum of a group isomorphic to $\mathbf{Q}/2\mathbf{Z}$ and a collection of groups isomorphic to \mathbf{Q} , one for each prime number p ; we write the group $\sqrt{\mathbf{Q}^*}$ multiplicatively, writing $(-1)^m$ for the element of $\sqrt{\mathbf{Q}^*}$ corresponding to $(m \bmod 2\mathbf{Z}) \in \mathbf{Q}/2\mathbf{Z}$, and p^m for the element of $\sqrt{\mathbf{Q}^*}$ corresponding to the element m of the p th summand \mathbf{Q} . The subgroup of $\sqrt{\mathbf{Q}^*}$ generated by $-1 (= (-1)^1)$ and all prime numbers $p (= p^1)$ identifies itself with \mathbf{Q}^* .

A set of coset representatives for $\sqrt{\mathbf{Q}^*}$ modulo \mathbf{Q}^* is given by the set \mathcal{R} of all elements of the form $(-1)^m \cdot \prod_{p \text{ prime}} p^{m(p)}$ where m and all $m(p)$ are rational numbers in the interval $[0, 1)$, almost all of them being equal to 0. Thus, by the proposition from the previous section, \mathcal{R} is a \mathbf{Q} -basis for the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$. That is, every element of $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ has a unique representation as a sum $\sum_{r \in \mathcal{R}} q_r \cdot r$, with $q_r \in \mathbf{Q}$ for all $r \in \mathcal{R}$ and $q_r = 0$ for almost all $r \in \mathcal{R}$. It is evident how to add and multiply two such expressions.

The field \mathbf{Q} does not satisfy Kneser's condition (v) in the theorem from Section 5, so (ii) fails as well, and the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ has zero-divisors. It is, indeed, not hard to write down zero-divisors.

For an odd prime number p , write ζ_p for the element $-(-1)^{1/p}$ of $\sqrt{\mathbf{Q}^*}$. It has order

p , so in the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ we have the identity

$$(1 - \zeta_p) \cdot (1 + \zeta_p + \zeta_p^2 + \dots + \zeta_p^{p-1}) = 0.$$

Also, since the elements $(-1)^{i/p}$ ($0 \leq i < p$) belong to \mathcal{R} , neither of the two factors vanishes in the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$, so they are both zero-divisors. Zero-divisors of this type correspond to the “harmless” additive relations between roots of unity that we encountered earlier. There are also zero-divisors reflecting the “entangled radicals”, as follows. For an odd prime number p , define the “Gauss sum” $\tau_p \in \mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ by

$$\tau_p = (-1)^{(p-1)/4} \cdot \sum_{i=1}^{p-1} \left(\frac{i}{p}\right) \cdot \zeta_p^i,$$

where the Legendre symbol $\left(\frac{i}{p}\right)$ is defined by $\left(\frac{i}{p}\right) \in \{1, -1\}$, $\left(\frac{i}{p}\right) \equiv i^{(p-1)/2} \pmod{p}$. Then one can show the identity $(p^{1/2} - \tau_p) \cdot (p^{1/2} + \tau_p) = 1 + \zeta_p + \dots + \zeta_p^{p-1}$, and therefore

$$(1 - \zeta_p) \cdot (p^{1/2} - \tau_p) \cdot (p^{1/2} + \tau_p) = 0 \quad \text{in } \mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}.$$

This identity is also valid for $p = 2$, if one defines $\zeta_2 = -1$ and $\tau_2 = (-1)^{1/4} - (-1)^{3/4}$.

There are no other zero-divisors that need concern us. To make this statement precise, recall that we are interested in turning the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ into a field containing $\sqrt{\mathbf{Q}^*}$ by factoring out by an ideal. For each p , we have a vanishing product of three elements of $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$. Any ring homomorphism from $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ to a field will necessarily map at least one of these three elements to 0; in fact, *exactly* one, since any two of the three factors $1 - \zeta_p$, $p^{1/2} - \tau_p$, $p^{1/2} + \tau_p$ are coprime in the sense of generating the unit ideal. We are only interested in maps to a field that are *injective* when restricted to $\sqrt{\mathbf{Q}^*}$, so none of the elements $1 - \zeta_p$ should map to 0. Therefore, for each prime number p , exactly one of $p^{1/2} - \tau_p$ and $p^{1/2} + \tau_p$ maps to zero. The following result expresses that this is the full story.

Proposition. Write \mathcal{P} for the set of prime numbers, and $\{1, -1\}^{\mathcal{P}}$ for the set of all maps $\mathcal{P} \rightarrow \{1, -1\}$. Also, denote by \mathcal{M} the set of ideals \mathfrak{m} of $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$ for which there exist an algebraically closed field Ω containing \mathbf{Q} and a ring homomorphism $\psi: \mathbf{Q}\{\sqrt{\mathbf{Q}^*}\} \rightarrow \Omega$ induced by an embedding $\sqrt{\mathbf{Q}^*} \rightarrow \Omega^*$, such that $\mathfrak{m} = \ker \psi$. Then there is a bijection $\{1, -1\}^{\mathcal{P}} \rightarrow \mathcal{M}$ that sends each map $\epsilon: \mathcal{P} \rightarrow \{1, -1\}$ to the ideal \mathfrak{m}_ϵ generated by $\{p^{1/2} - \epsilon(p)\tau_p : p \text{ prime}\}$.

The proposition implies that all of the rings $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}/\mathfrak{m}_\epsilon$ are fields. All these fields are isomorphic to $\mathbf{Q}(\sqrt{\mathbf{Q}^*})$, and therefore also to each other, although no two of the ideals \mathfrak{m}_ϵ are equal.

Using the proposition, one can construct a basis for $\mathbf{Q}(\sqrt{\mathbf{Q}^*})$ as a vector space over \mathbf{Q} . Let \mathcal{R}' consist of all elements of the form $\prod_{p \text{ prime}} ((-1)^{2l(p)} \cdot p^{m(p)})$, where $l(p) \in \mathbf{Z}[1/p]$ and $m(p) \in \mathbf{Q}$ satisfy $0 \leq l(p) < 1 - 1/p$, $0 \leq m(p) < 1/2$, and $l(p) = m(p) = 0$ for almost all p . Then for each $\epsilon \in \{1, -1\}^{\mathcal{P}}$, the natural map $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\} \rightarrow \mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}/\mathfrak{m}_\epsilon$ sends \mathcal{R}' bijectively to a basis for $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}/\mathfrak{m}_\epsilon$ over \mathbf{Q} . Likewise, \mathcal{R}' yields a basis for $\mathbf{Q}(\sqrt{\mathbf{Q}^*})$ over \mathbf{Q} .

Again we can bring in the Galois group $\text{Gal}(K(\sqrt{K^*})/K)$. The set $\{1, -1\}^{\mathcal{P}}$ has an obvious group structure, and it turns out that $\text{Gal}(\mathbf{Q}(\sqrt{\mathbf{Q}^*})/\mathbf{Q})$ may be identified with the kernel of a naturally defined surjective group homomorphism $\text{Aut}_{\mathbf{Q}^*} \sqrt{\mathbf{Q}^*} \rightarrow \{1, -1\}^{\mathcal{P}}$.

The conclusion is that the case $K = \mathbf{Q}$ is well under control. The extension to general K requires some Galois-theoretic preparation. One of the conclusions will be that for any K there is an abelian group E_K that plays the role that $\{1, -1\}^{\mathcal{P}}$ plays in the case $K = \mathbf{Q}$.

7. Profinite integers

Let A be an abelian group. An *endomorphism* of A is a group homomorphism $A \rightarrow A$. Writing A multiplicatively, we can define a ring structure on the set $\text{End } A$ of all endomorphisms of A by putting $(f + g)(a) = f(a) \cdot g(a)$ and $(fg)(a) = f(g(a))$ for $f, g \in \text{End } A$, $a \in A$. The unit element of $\text{End } A$ is the identity map $A \rightarrow A$, and the group $(\text{End } A)^*$ of units of $\text{End } A$ equals the automorphism group $\text{Aut } A$ of A .

For example, if n is a positive integer and A is cyclic of order n , then there is a ring isomorphism $\mathbf{Z}/n\mathbf{Z} \rightarrow \text{End } A$ that sends $(i \bmod n)$ to the endomorphism $a \mapsto a^i$ of A , for each $i \in \mathbf{Z}$.

Let Ω be an algebraically closed field of characteristic zero. For a positive integer n , we write μ_n for the group $\Omega^*[n]$ of n th roots of unity in Ω ; it is cyclic of order n . Also, we write $\mu = \bigcup_{n \geq 1} \mu_n$, which is the group of all roots of unity in Ω .

The ring $\hat{\mathbf{Z}}$ of *profinite integers* may for our purposes be defined by $\hat{\mathbf{Z}} = \text{End } \mu$. Since each endomorphism of μ sends each subgroup μ_n to itself, $\hat{\mathbf{Z}}$ may be considered as a subring of the product ring $\prod_{n \geq 1} \text{End } \mu_n \cong \prod_{n \geq 1} \mathbf{Z}/n\mathbf{Z}$ (with component-wise ring operations), and it is not hard to figure out which elements $(a_n)_{n=1}^\infty$ of $\prod_{n \geq 1} \mathbf{Z}/n\mathbf{Z}$ correspond to elements of $\hat{\mathbf{Z}}$: exactly those that satisfy $a_m \equiv a_n \pmod{m}$ for any two positive integers m, n with m dividing n . This description, which is in fact the more usual definition of $\hat{\mathbf{Z}}$,

shows that $\hat{\mathbf{Z}}$ is independent of the choice of Ω . It also shows that $\hat{\mathbf{Z}}$ is a *commutative* ring, which does not hold for all endomorphism rings of abelian groups.

The group $\hat{\mathbf{Z}}^*$ of units of $\hat{\mathbf{Z}}$ equals the automorphism group $\text{Aut } \mu$ of μ . It may be identified with the multiplicative group consisting of those elements $(a_n)_{n=1}^\infty$ of $\prod_{n \geq 1} \mathbf{Z}/n\mathbf{Z}$ that, in addition to satisfying the system of congruences $a_m \equiv a_n \pmod{m}$ just considered, have the property $a_n \in (\mathbf{Z}/n\mathbf{Z})^*$ for all n .

One cannot do infinite Galois theory without understanding the basic properties of the ring $\hat{\mathbf{Z}}$. A recreational introduction to $\hat{\mathbf{Z}}$ and some of its properties is found in the paper [21] from which the illustration on the title page is taken; for more formal treatments, one may consult [25, 32]. For our purposes, the following will suffice.

Let p be a prime number. The ring \mathbf{Z}_p of *p-adic integers* is defined to be the endomorphism ring of the subgroup $\bigcup_{i=0}^\infty \mu_{p^i}$ of μ . The ring \mathbf{Z}_p admits a description similar to the one we gave of $\hat{\mathbf{Z}}$, but with the set of positive integers n replaced by the set of powers of p . Other than $\hat{\mathbf{Z}}$, it is a *domain*.

As an infinite analogue of the Chinese remainder theorem, one has an isomorphism $\hat{\mathbf{Z}} \cong \prod_{p \text{ prime}} \mathbf{Z}_p$ of rings, as well as an induced group isomorphism $\hat{\mathbf{Z}}^* \cong \prod_{p \text{ prime}} \mathbf{Z}_p^*$.

8. Profinite groups

A *directed partially ordered set* is a set I with a partial ordering \leq such that for any $i, j \in I$ there exists $k \in I$ with $i \leq k$ and $j \leq k$. A *projective system* of groups is a directed partially ordered set I , together with a collection $(G_i)_{i \in I}$ of groups and a collection of group homomorphisms $(\rho_i^j: G_j \rightarrow G_i)_{i, j \in I, i \leq j}$, such that ρ_i^i is the identity on G_i for each i , and $\rho_i^k = \rho_i^j \circ \rho_j^k$ whenever $i, j, k \in I$ satisfy $i \leq j \leq k$. The *projective limit* $\varprojlim G_i$ of such a system is defined to be the subgroup of $\prod_{i \in I} G_i$ consisting of all elements $(g_i)_{i \in I}$ satisfying $\rho_i^j(g_j) = g_i$ whenever $i, j \in I, i \leq j$. Equipping each G_i with the discrete topology, $\prod_{i \in I} G_i$ with the product topology, and the projective limit with the restriction topology, we find that $\varprojlim G_i$ is a *topological group*.

A *profinite group* is a topological group that is isomorphic to the projective limit of a projective system of *finite* groups G_i ; or, equivalently, that is totally disconnected and compact. For more information on profinite groups, one may consult [6, Chapter V; 25; 32].

The additive group of the ring $\hat{\mathbf{Z}}$ is a profinite group, and so is its group $\hat{\mathbf{Z}}^*$ of units. In these cases, I is the set of positive integers, ordered by divisibility. With the same I , we define the group $T(\mathbf{G}_m)$, the *Tate module of the multiplicative group*, in the following

way. For a positive integer n , let, as in the previous section, μ_n be the group of n th roots of unity, and for m dividing n let the map $\mu_n \rightarrow \mu_m$ map each ζ to $\zeta^{n/m}$; this defines a projective system, and $T(\mathbf{G}_m)$ is its projective limit. The reader may verify that $T(\mathbf{G}_m)$ is in fact isomorphic to $\hat{\mathbf{Z}}$; more canonically, $T(\mathbf{G}_m)$ is a free $\hat{\mathbf{Z}}$ -module of rank 1.

The group $\text{Aut}_{K^*} \sqrt{K^*}$ that we encountered in the introduction, is also an example of a profinite group: if one takes I to be the set of all subgroups $A \subset \sqrt{K^*}$ for which there exists a positive integer n with $\#A[n] = n$, $A^n \subset K^* \subset A$, and $(A : K^*) < \infty$, ordered by inclusion, then $\text{Aut}_{K^*} A$ is finite for each $A \in I$, and one has $\text{Aut}_{K^*} \sqrt{K^*} = \varprojlim \text{Aut}_{K^*} A$.

The group $\hat{\mathbf{Z}}$ plays among the profinite groups the same role that \mathbf{Z} plays among all groups. More specifically, for each profinite group G and each $\gamma \in G$, there is a unique group homomorphism $\hat{\mathbf{Z}} \rightarrow G$ that sends 1 to γ , and it is automatically continuous; the image of $a \in \hat{\mathbf{Z}}$ under this group homomorphism is written γ^a .

9. Infinite Galois theory

Infinite Galois theory studies Galois extensions that are not required to be finite, such as the extension $K \subset K(\sqrt{K^*})$. The definition is as follows.

Let $K \subset L$ be an extension of fields, and write $\mathcal{U} = \{M : M \text{ is a subfield of } L \text{ containing } K, \text{ and } M \text{ is a finite Galois extension of } K\}$. The set \mathcal{U} , ordered by inclusion, is a directed partially ordered set. We say that L is a *Galois extension* of K if L is equal to the union of all $M \in \mathcal{U}$.

Suppose $K \subset L$ is a Galois extension. As in the finite case, the *Galois group* $\text{Gal}(L/K)$ is defined to be the group of all field automorphisms of L that are the identity on K . With \mathcal{U} as above, one has $\sigma M = M$ for each $\sigma \in \text{Gal}(L/K)$ and each $M \in \mathcal{U}$, and this leads to restriction maps $\text{Gal}(L/K) \rightarrow \text{Gal}(M/K)$ (for $M \in \mathcal{U}$) and to a group isomorphism $\text{Gal}(L/K) \cong \varprojlim \text{Gal}(M/K)$, the projective limit being taken over \mathcal{U} . This isomorphism endows $\text{Gal}(L/K)$ with a topology, the *Krull topology*; a subset $U \subset \text{Gal}(L/K)$ is open if and only if for each $\sigma \in U$ there exists $\alpha \in L$ with the property $\{\tau \in \text{Gal}(L/K) : \tau(\alpha) = \sigma(\alpha)\} \subset U$.

All the usual theorems from finite Galois theory have an extension to infinite Galois theory. For example, a field extension L of K is Galois if and only if one can obtain L by adjoining to K , inside a suitable algebraically closed field, all zeroes of a collection of polynomials $f \in K[X]$ that are *separable* in the sense that f and $\frac{df}{dX}$ generate the unit ideal of $K[X]$. The well-known bijection, from finite Galois theory, between the set of

intermediate fields and the set of subgroups of the Galois group, remains valid provided that one restricts to the set of *closed* subgroups of $\text{Gal}(L/K)$, as usual in topological algebra (see [6, Chapter V, Theorem 2]).

We give three types of examples of Galois extensions and their groups. First, suppose K is a *finite* field, and let L be an algebraic closure of K . Then L is a Galois extension of K , and its Galois group contains the “Frobenius” element F that maps each $x \in L$ to $x^{\#K}$. Also, there is an isomorphism $\hat{\mathbf{Z}} \cong \text{Gal}(L/K)$ that maps each $a \in \hat{\mathbf{Z}}$ to F^a .

Secondly, let K be any field of characteristic zero, and let μ be the group of all roots of unity in some algebraically closed field that contains K . Then $K(\mu)$ is a Galois extension of K . Restricting the action of the Galois group to μ , we obtain a group homomorphism $\text{Gal}(K(\mu)/K) \rightarrow \text{Aut } \mu = \hat{\mathbf{Z}}^*$. It induces an isomorphism of topological groups between $\text{Gal}(K(\mu)/K)$ and a certain closed subgroup Γ_K of $\hat{\mathbf{Z}}^*$, the *cyclotomic Galois group* of K . For example, one has $\Gamma_{\mathbf{Q}} = \hat{\mathbf{Z}}^*$, by the irreducibility of the cyclotomic polynomials over \mathbf{Q} .

Finally, let K be any field of characteristic zero, and let $\sqrt{K^*}$ be embedded in the multiplicative group of an algebraically closed field containing K . Then the extension $K \subset K(\sqrt{K^*})$ is a Galois extension, and in Section 4 we defined a restriction map $\text{Gal}(K(\sqrt{K^*})/K) \rightarrow \text{Aut}_{K^*} \sqrt{K^*}$ that enables us to view $\text{Gal}(K(\sqrt{K^*})/K)$ as a subgroup of $\text{Aut}_{K^*} \sqrt{K^*}$. Both $\text{Gal}(K(\sqrt{K^*})/K)$ and $\text{Aut}_{K^*} \sqrt{K^*}$ are profinite groups, and one verifies that the restriction map is continuous; it follows that $\text{Gal}(K(\sqrt{K^*})/K)$ may be viewed as a *closed* subgroup of $\text{Aut}_{K^*} \sqrt{K^*}$ and that the Krull topology coincides with the restriction topology. In Section 13 we address the question *which* subgroup it is.

10. Galois algebras

It is usually much easier to recognize that an algebraic object is a ring than to tell whether it is a field. For example, $\mathbf{Z}/n\mathbf{Z}$ is a ring for each integer n , but the question whether it is a field boils down to primality testing, a fairly subtle subject that depends on ring theory in the first place. Similarly, a *Galois algebra* over a field “looks like” a Galois extension, except that it is only a ring and may or may not be a field.

Here are the definitions. Let K be a field. In these lectures, a *K-algebra* is a commutative ring F equipped with a ring homomorphism $K \rightarrow F$, and if F, F' are K -algebras, then a *K-algebra homomorphism* $F \rightarrow F'$ is a ring homomorphism such that the composed map $K \rightarrow F \rightarrow F'$ coincides with the given map $K \rightarrow F'$. Next let G be a profinite group. A *Galois algebra over K with group G* is a non-zero K -algebra F equipped with an action of G on F by K -algebra automorphisms such that the following three conditions are satisfied:

- (i) the action $G \times F \rightarrow F$ is *continuous* if F is given the discrete topology, or, equivalently, for each $a \in F$ the subgroup $\{\sigma \in G : \sigma a = a\}$ of G is *open*;
- (ii) if Ω is an algebraically closed field containing K , then for any two K -algebra homomorphisms $s, t: F \rightarrow \Omega$ there is a unique element $\sigma \in G$ such that $s = t \circ \sigma$;
- (iii) each element of F is a zero of a separable polynomial in $K[X]$ (it suffices to require this for all elements in a set of K -algebra generators for F).

A few examples may be more enlightening than the definition. Any Galois extension $K \subset L$ of fields is a Galois algebra with group $\text{Gal}(L/K)$. If L is a Galois extension of \mathbf{Q} with ring of integers R , and p is a prime number that does not ramify in R , then R/pR is a Galois algebra over $\mathbf{Z}/p\mathbf{Z}$ with group $\text{Gal}(L/\mathbf{Q})$; it is not necessarily a field. For any field K and each integer $n > 1$, the K -algebra $F = K \times K \times \dots \times K$, with n factors K and componentwise operations, is an example of a Galois algebra that is not a field, with a cyclic group $G = \langle \sigma \rangle$ of order n that acts on F by

$$\sigma(a_0, a_1, \dots, a_{n-1}) = (a_1, a_2, \dots, a_0), \quad \text{for } a_0, a_1, \dots, a_{n-1} \in K.$$

More generally, if G is any profinite group, and K is a field, then the K -algebra of all continuous functions $h: G \rightarrow K$ (where K is discrete) contains K in a natural way, it is acted upon by G by $(\tau h)(\sigma) = h(\sigma\tau)$ (for $\sigma, \tau \in G$), and it is a Galois algebra over K with group G ; an algebra of this type is called *totally split*.

A slight generalization of the last example given, comprises in fact *all* Galois algebras. Suppose that $K \subset L$ is a Galois extension of fields, and that its Galois group $\text{Gal}(L/K)$ is embedded as a closed subgroup in a profinite group G . Consider the K -algebra of all continuous maps $h: G \rightarrow L$ that respect the $\text{Gal}(L/K)$ -action in the sense that $h(\delta\sigma) = \delta(h(\sigma))$ for all $\delta \in \text{Gal}(L/K)$, $\sigma \in G$. It is acted upon by G by $(\tau h)(\sigma) = h(\sigma\tau)$ for all $\sigma, \tau \in G$. With this action, it is a Galois algebra over K with group G . If $\text{Gal}(L/K)$ has finite index in G , then this algebra of functions is just the product of $(G : \text{Gal}(L/K))$ fields isomorphic to L .

To see that any Galois algebra F over a field K with group G is of this form, one considers its *spectrum* $\text{Spec } F$. It is by definition the set of prime ideals of F . It is non-empty, and coincides in fact with the set of *maximal* ideals of F . One proves that the obvious action of G on $\text{Spec } F$ is *transitive*. Fixing $\mathfrak{m} \in \text{Spec } F$, one defines the *decomposition* group $D = D_{\mathfrak{m}}$ by $D = \{\sigma \in G : \sigma\mathfrak{m} = \mathfrak{m}\}$, so that there is a bijection $G/D \rightarrow \text{Spec } F$ sending the coset σD to the image $\sigma\mathfrak{m}$ of \mathfrak{m} under σ . The decomposition group is a closed subgroup

of G , and a different choice of \mathfrak{m} replaces D by a conjugate subgroup of G . Now one proves that the field F/\mathfrak{m} is a Galois extension of K with Galois group equal to D , and that F itself may be identified with the K -algebra consisting of all continuous maps $h: G \rightarrow F/\mathfrak{m}$ that respect the D -action; this identification respects the G -action.

The overall conclusion is that any Galois algebra is “fairly close” to being a field, and that the decomposition group $D \subset G$ measures exactly “how close”. For example, the Galois algebra is a field if and only if one has $D = G$, and it is totally split if and only if D is trivial.

If one is faced with a Galois algebra, the first thing to do is determine the decomposition group. It is a closed subgroup of G that is well-defined up to conjugation.

11. The ring of radicals

Having prepared our way, we return to the problem we were studying. Let K be a field of characteristic zero, let $\sqrt{K^*}$ be a group of radicals of K^* as in Section 4, and let $K\{\sqrt{K^*}\}$ be the ring from Section 5; it contains K as a subring and $\sqrt{K^*}$ as a subgroup of its group $K\{\sqrt{K^*}\}^*$ of units, and it was defined in such a way that the two induced inclusions $K^* \subset K\{\sqrt{K^*}\}^*$ are the same. We are interested in understanding the ideals \mathfrak{m} of $K\{\sqrt{K^*}\}$ for which there exist an algebraically closed field Ω containing K and a K -algebra homomorphism $\psi: K\{\sqrt{K^*}\} \rightarrow \Omega$ induced by an embedding $\sqrt{K^*} \rightarrow \Omega^*$ (as defined in Section 4), such that $\mathfrak{m} = \ker \psi$. All such \mathfrak{m} are maximal ideals.

We start by taking care of what we called the “harmless” relations between roots of unity. Let $C \subset \sqrt{K^*}$ be a subgroup of prime order that is not contained in K^* . Then for each $\zeta \in C$, $\zeta \neq 1$ one has

$$(1 - \zeta) \cdot \sum_{\eta \in C} \eta = 0, \quad 1 - \zeta \neq 0, \quad \sum_{\eta \in C} \eta \neq 0$$

in $K\{\sqrt{K^*}\}$. As in Section 6, the ring homomorphisms $K\{\sqrt{K^*}\} \rightarrow \Omega$ that we are interested in do not map $1 - \zeta$ to 0. Hence, they map $\sum_{\eta \in C} \eta$ to zero, so any of the ideals that we are interested in contains $\sum_{\eta \in C} \eta$. Thus, we may as well factor out $K\{\sqrt{K^*}\}$ by the ideal generated by all elements obtained in this way. That is, we replace $K\{\sqrt{K^*}\}$ by the ring $K\langle\sqrt{K^*}\rangle$ defined by

$$K\langle\sqrt{K^*}\rangle = K\{\sqrt{K^*}\}/\mathfrak{b},$$

where \mathfrak{b} is the ideal of $K\{\sqrt{K^*}\}$ generated by $\{\sum_{\eta \in C} \eta : C \text{ is a subgroup of prime order of } \sqrt{K^*} \text{ not contained in } K^*\}$.

Just like $K\{\sqrt{K^*}\}$, the ring $K\langle\sqrt{K^*}\rangle$ contains K as a subring and $\sqrt{K^*}$ as a subgroup of its group of units. It also has two noteworthy properties that $K\{\sqrt{K^*}\}$ fails to have (unless $\mathfrak{b} = 0$).

The first property is that *any* maximal ideal of $K\langle\sqrt{K^*}\rangle$ occurs as the kernel of a K -algebra homomorphism from $K\langle\sqrt{K^*}\rangle$ to an algebraically closed field Ω containing K induced by an embedding $\sqrt{K^*} \rightarrow \Omega^*$; equivalently, the set of maximal ideals described at the beginning of the present section equals the closed subset of $\text{Spec } K\{\sqrt{K^*}\}$ defined by \mathfrak{b} . In particular, the set \mathcal{M} considered in the proposition of Section 6 may be identified with $\text{Spec } \mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$. This first property follows from the alternative description

$$K\langle\sqrt{K^*}\rangle = S^{-1}K\{\sqrt{K^*}\},$$

where S is the multiplicative subset of $K\{\sqrt{K^*}\}$ generated by $\{a - 1 : a \in \sqrt{K^*}, a \neq 1\}$.

The second property of $K\langle\sqrt{K^*}\rangle$ is expressed in the following result. Due to the canonical nature of the construction of the K -algebra $K\langle\sqrt{K^*}\rangle$ out of the group $\sqrt{K^*}$, the profinite group $\text{Aut}_{K^*}\sqrt{K^*}$ acts in a natural way as a group of K -algebra automorphisms on $K\langle\sqrt{K^*}\rangle$, and the action is continuous. Remember also that $\text{Gal}(K(\sqrt{K^*})/K)$ may be viewed as a subgroup of $\text{Aut}_{K^*}\sqrt{K^*}$; this subgroup is, just as the decomposition group in the previous section, only well-defined up to conjugation.

Theorem. *The K -algebra $K\langle\sqrt{K^*}\rangle$ is a Galois algebra over K with group $\text{Aut}_{K^*}\sqrt{K^*}$, and the decomposition group is conjugate to $\text{Gal}(K(\sqrt{K^*})/K)$.*

The proof makes use of the first property that we formulated.

We can make sure that the subgroup $\text{Gal}(K(\sqrt{K^*})/K)$ and the decomposition group $D = D_{\mathfrak{m}}$ coincide exactly, not just up to conjugation. To achieve this, it suffices to let \mathfrak{m} be the kernel of the K -algebra homomorphism corresponding to the embedding that we use to view $\text{Gal}(K(\sqrt{K^*})/K)$ as a subgroup of $\text{Aut}_{K^*}\sqrt{K^*}$. In Section 13 we shall see that this precaution is actually unnecessary.

From the previous section, one now sees that $K\langle\sqrt{K^*}\rangle$ is a field if and only if one has $\text{Gal}(K(\sqrt{K^*})/K) = \text{Aut}_{K^*}\sqrt{K^*}$. By the theorem of Section 5, this occurs if and only if $K\{\sqrt{K^*}\}$ is a field.

The importance of the ring $K\langle\sqrt{K^*}\rangle$ merits that it gets a special name; we shall call it the *ring of radicals over K* .

12. The algebra of the ring of radicals

We continue to assume that K is a field of characteristic zero. We may summarize the definition of the ring $K\langle\sqrt{K^*}\rangle$ of radicals over K by saying that, as a K -algebra, it is generated by the set $\sqrt{K^*}$ subject to three types of relations: the multiplicative relations in the group $\sqrt{K^*}$; the relations saying that each $a \in K^*$ is the same when viewed as an element of K or as an element of $\sqrt{K^*}$; and one additional relation $\sum_{\eta \in \sqrt{K^*}[p]} \eta = 0$ for each prime number p with $\#K^*[p] = 1$.

In the present section we write down a basis for $K\langle\sqrt{K^*}\rangle$ as a vector space over K , and we briefly discuss its spectrum.

We start by considering the case $K = \mathbf{Q}$, using the notation from Section 6. Consider the collection of elements of $\sqrt{\mathbf{Q}^*}$ of the form $\prod_{p \text{ prime}} (-1)^{2l(p)} \cdot p^{m(p)}$, where $l(p) \in \mathbf{Z}[1/p]$ and $m(p) \in \mathbf{Q}$ satisfy $0 \leq l(p) < 1$, $l(2) < 1/2$, $0 \leq m(p) < 1$ for all p , and $l(p) = m(p) = 0$ for almost all p . Up to signs, this collection coincides with the set \mathcal{R} defined in Section 6, so just like \mathcal{R} it provides a basis for the ring $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ over \mathbf{Q} . Now restrict to the subset \mathcal{S} of this collection that is defined by the inequalities $l(p) < 1 - 1/p$ for all p . Then \mathcal{S} provides a basis for $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ over \mathbf{Q} . The “missing” elements $(-1)^{2l(p)}$, for p odd and $1 - 1/p \leq l(p) < 1$, are expressed in \mathcal{S} through $(-1)^{2l(p)} = \sum_{i=1}^{p-1} (-1)^{2(l(p)-i/p)}$; this identity in $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ is a consequence of the relation $\sum_{j=0}^{p-1} (-1)^{2j/p} = \sum_{\eta \in \sqrt{\mathbf{Q}^*}[p]} \eta = 0$. It is now clear how to add and multiply two elements of $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ that are both written as \mathbf{Q} -linear combinations of elements of \mathcal{S} .

For general K , we write E for a divisible hull of K^* , and we choose a group isomorphism $\sqrt{K^*} \cong E \oplus \bigoplus_p (\mathbf{Z}[1/p]/\mathbf{Z})$ sending each $b \in K^*$ to $(b, 0)$, with p ranging over the set of prime numbers with $\#K^*[p] = 1$ (cf. Section 4). Let \mathcal{T}' be a set of coset representatives for E modulo K^* , and let $\mathcal{T} \subset \sqrt{K^*}$ correspond to the subset $\{(s, (s_p \bmod \mathbf{Z})_p) : s \in \mathcal{T}', s_p \in \mathbf{Z}[1/p], 0 \leq s_p < 1 - 1/p \text{ for all } p, \text{ and } s_p = 0 \text{ for almost all } p\}$ of $E \oplus \bigoplus_p (\mathbf{Z}[1/p]/\mathbf{Z})$. The proposition from Section 5, which gives a basis for $K\langle\sqrt{K^*}\rangle$ over K , now has the following analogue for $K\langle\sqrt{K^*}\rangle$.

Proposition. *The set \mathcal{T} forms a basis for $K\langle\sqrt{K^*}\rangle$ as a vector space over K .*

Again, the reader can easily figure out a rule for multiplying two elements of $K\langle\sqrt{K^*}\rangle$ represented on this basis, so the arithmetic in $K\langle\sqrt{K^*}\rangle$ is well under control when K and E are taken for granted.

Next we discuss the spectrum $\text{Spec } K\langle\sqrt{K^*}\rangle$ of the ring of radicals. In the case $K = \mathbf{Q}$, this spectrum may be identified with the set \mathcal{M} that occurs in the proposition from

Section 6. Thus, we see from that proposition that $\text{Spec } \mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ is in bijection with the set $\{1, -1\}^{\mathcal{P}}$, where \mathcal{P} is the set of all prime numbers. For general K , we know from the previous section that $K\langle\sqrt{K^*}\rangle$ is a Galois algebra over K with group $\text{Aut}_{K^*}\sqrt{K^*}$ and decomposition group $\text{Gal}(K(\sqrt{K^*})/K)$, so by Section 10 its spectrum is in bijection with the coset space $\text{Aut}_{K^*}\sqrt{K^*}/\text{Gal}(K(\sqrt{K^*})/K)$.

This discussion shows that for $K = \mathbf{Q}$ the coset space $\text{Aut}_{K^*}\sqrt{K^*}/\text{Gal}(K(\sqrt{K^*})/K)$ may be identified with the set $\{1, -1\}^{\mathcal{P}}$. One may wonder what it looks like for general K . We address this question in the following section.

13. The radical Galois group

Let K be a field of characteristic zero, and denote by μ the multiplicative group of all roots of unity in an algebraically closed field containing K . In Section 9, we saw that $\text{Gal}(K(\mu)/K)$ may be identified with a certain closed subgroup Γ_K of $\hat{\mathbf{Z}}^*$. Write W_K for the closure of the $\hat{\mathbf{Z}}$ -ideal generated by $\{\gamma - 1 : \gamma \in \Gamma_K\}$. If the number of roots of unity in K^* is finite, say w_K , then one has $W_K = \hat{\mathbf{Z}}w_K$, and in general W_K is the annihilator of the subgroup $\mu \cap K^*$ of μ in the ring $\hat{\mathbf{Z}} = \text{End } \mu$. One may say that the ideal W_K measures “how many” roots of unity there are in K^* .

One can now prove that the set $\Gamma_K^{W_K} = \{\gamma^a : \gamma \in \Gamma_K, a \in W_K\}$ is a closed subgroup of the multiplicative group $\hat{\mathbf{Z}}^* \cap (1 + W_K^2)$ of all units of $\hat{\mathbf{Z}}$ that are 1 modulo the square of the ideal W_K . We define

$$E_K = (\hat{\mathbf{Z}}^* \cap (1 + W_K^2))/\Gamma_K^{W_K}.$$

This is a profinite abelian group, the *entanglement group* of K .

The field \mathbf{Q} of rational numbers may again serve as an example. As we saw in Section 9, we have $\Gamma_{\mathbf{Q}} = \hat{\mathbf{Z}}^*$, and in Section 7 we saw that this group may be identified with $\prod_p \text{prime } \mathbf{Z}_p^*$. Also, since the number of roots of unity in \mathbf{Q}^* equals 2, we have $W_{\mathbf{Q}} = 2\hat{\mathbf{Z}}$. The group $\hat{\mathbf{Z}}^* \cap (1 + W_{\mathbf{Q}}^2) = \hat{\mathbf{Z}}^* \cap (1 + 4\hat{\mathbf{Z}})$ may now be identified with the group $(1 + 4\mathbf{Z}_2) \times \prod_{p>2 \text{ prime}} \mathbf{Z}_p^*$, whereas $\Gamma_{\mathbf{Q}}^{W_{\mathbf{Q}}} = \hat{\mathbf{Z}}^{*2}$ may be identified with $(1 + 8\mathbf{Z}_2) \times \prod_{p>2 \text{ prime}} \mathbf{Z}_p^{*2}$. Since each of the groups $(1 + 4\mathbf{Z}_2)/(1 + 8\mathbf{Z}_2)$ and $\mathbf{Z}_p^*/\mathbf{Z}_p^{*2}$ (with p an odd prime) is of order 2, we find a canonical isomorphism

$$E_{\mathbf{Q}} = (\hat{\mathbf{Z}}^* \cap (1 + 4\hat{\mathbf{Z}}))/\hat{\mathbf{Z}}^{*2} \cong \{1, -1\}^{\mathcal{P}},$$

where $\{1, -1\}^{\mathcal{P}}$ is the group that we encountered in Section 6.

As an exercise in understanding the group $\hat{\mathbf{Z}}^*$, the reader may prove that the group E_K is trivial if and only if K^* has an element of order 4 as well as, for each odd prime number p , an element of order p .

We can now answer the question posed at the end of the previous section.

Theorem. *There is a continuous group homomorphism $\text{Aut}_{K^*}\sqrt{K^*} \rightarrow E_K$ such that the following is true. Let Ω be an algebraically closed field containing K , let $\sqrt{K^*} \rightarrow \Omega^*$ be an embedding, and let $\text{Gal}(K(\sqrt{K^*})/K) \rightarrow \text{Aut}_{K^*}\sqrt{K^*}$ be the induced group homomorphism. Then the sequence*

$$1 \rightarrow \text{Gal}(K(\sqrt{K^*})/K) \rightarrow \text{Aut}_{K^*}\sqrt{K^*} \rightarrow E_K \rightarrow 1$$

is an exact sequence of not necessarily abelian profinite groups.

This beautiful theorem was proved by B. de Smit and W. J. Palenstijn [10], who elaborated upon work done by M. Honsbeek [14, Chapter 1]. The proof will be given in the special session accompanying the present series of lectures. It depends on Schinzel’s theorem on “abelian binomials” [26, Theorem 2; 30, Chapitre IV, Exercice 1.4].

One can also say *which* map $\text{Aut}_{K^*}\sqrt{K^*} \rightarrow E_K$ has the property described in the theorem. This is important to know for the applications.

The entanglement group measures the extent to which entanglement of radicals occurs. It is trivial if and only if one has

$$K\{\sqrt{K^*}\} = K\langle\sqrt{K^*}\rangle \cong K(\sqrt{K^*}).$$

From the definition we see that E_K depends only on the cyclotomic Galois group Γ_K of K . Thus, without changing E_K , we may replace K by the subfield $K' = K \cap \mathbf{Q}(\mu)$, which has the same cyclotomic Galois group. The group Γ_K acts on $\mathbf{Q}(\mu)$, and K' equals the field of invariants. All entangled radicals over K arise from entangled radicals over K' . Since all elements of K' are algebraic numbers, it may be said that entanglement of radicals is ultimately a number-theoretic phenomenon.

The theorem implies a remarkable fact: $\text{Gal}(K(\sqrt{K^*})/K)$ is a *normal* subgroup of $\text{Aut}_{K^*}\sqrt{K^*}$. Thus, it is independent of the choice of the embedding $\sqrt{K^*} \rightarrow \Omega^*$. Put in another way: all $\mathfrak{m} \in \text{Spec } K\langle\sqrt{K^*}\rangle$ have the same decomposition group $D_{\mathfrak{m}}$ in $\text{Aut}_{K^*}\sqrt{K^*}$. In the following section we shall see how this property can be interpreted.

Extending the theory developed so far to other situations would be of great interest. The case of non-zero characteristic p should be fairly immediate; we did not include it

in our discussion in order to avoid being distracted by inseparability phenomena and the absence of p th roots of unity. One may also replace the multiplicative group \mathbf{G}_m by other commutative algebraic groups, such as tori or elliptic curves, so that the role of the “radicals” is played by “division points” on these groups. In all these cases, the corresponding Galois group deserves being investigated.

14. The decomposition algebra

For a field K of characteristic zero, the K -algebra $K\langle\sqrt{K^*}\rangle$ is a Galois algebra over K with group $\text{Aut}_{K^*}\sqrt{K^*}$. For any $\mathfrak{m} \in \text{Spec } K\langle\sqrt{K^*}\rangle$, the field $K\langle\sqrt{K^*}\rangle/\mathfrak{m}$ is isomorphic to the field $K(\sqrt{K^*})$ that we are interested in. Thus, the question poses itself how to construct such an \mathfrak{m} . We do know $D_{\mathfrak{m}}$, from the theorems in Sections 11 and 13.

Assume, generally, that K is any field, that G is a profinite group, and that F is a Galois algebra over K with group G . Fix $\mathfrak{m} \in \text{Spec } F$, and let $D = D_{\mathfrak{m}}$ be its decomposition group. From Section 10 we know that the field F/\mathfrak{m} is a Galois extension of K with group D . Write $F^D = \{a \in F : \delta a = a \text{ for all } \delta \in D\}$, which is a sub- K -algebra of F , the *decomposition algebra* of \mathfrak{m} . The image of F^D under the map $F \rightarrow F/\mathfrak{m}$ equals the subfield K of F/\mathfrak{m} , so that we obtain a map $F^D \rightarrow K$. Now one can show that the natural map $F \otimes_{F^D} K \rightarrow F/\mathfrak{m}$ is an isomorphism of K -algebras or, equivalently, that the F -ideal \mathfrak{m} is generated by the kernel of $F^D \rightarrow K$.

Thus, both \mathfrak{m} and F/\mathfrak{m} may be described in terms of a certain K -algebra homomorphism $F^D \rightarrow K$. Assume now that D is a *normal* subgroup of G , which is the case in our motivating example. Then D and F^D are independent of the choice of \mathfrak{m} , and we call F^D the *decomposition algebra* of F rather than of \mathfrak{m} . It is a *totally split Galois algebra* over K with group G/D . This implies that there are “just as many” K -algebra homomorphisms $F^D \rightarrow K$ as there are elements in G/D ; that is, if $\phi: F^D \rightarrow K$ is one of them, then each such homomorphism is of the form $\phi \circ \sigma$ for a unique element $\sigma \in G/D$.

We return to the situation where K has characteristic zero, G equals $\text{Aut}_{K^*}\sqrt{K^*}$, and F is $K\langle\sqrt{K^*}\rangle$. By the theorem of Section 13, the decomposition group D is the kernel of a surjective map $G \rightarrow \mathbf{E}_K$, so G/D identifies itself with \mathbf{E}_K . Since the map $G \rightarrow \mathbf{E}_K$ can be made explicit, it is reasonable to expect that the decomposition algebra $K\langle\sqrt{K^*}\rangle^D$ can also be written down explicitly. It is a Galois algebra over K with an abelian group \mathbf{E}_K , and it admits a plentiful supply of K -algebra homomorphisms to K . For any one of them, the K -algebra

$$K\langle\sqrt{K^*}\rangle \otimes_{K\langle\sqrt{K^*}\rangle^D} K$$

may serve as an “explicit model” for the field $K(\sqrt{K^*})$.

As an example, we consider the case $K = \mathbf{Q}$. In Section 6 we defined, for each prime number p , elements $p^{1/2}$ and τ_p of the ring $\mathbf{Q}\{\sqrt{\mathbf{Q}^*}\}$. We denote their images in $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ also by $p^{1/2}$ and τ_p . In the latter ring, each of $p^{1/2}$ and τ_p has square equal to p . The decomposition algebra of $\mathbf{Q}\langle\sqrt{\mathbf{Q}^*}\rangle$ is, as a \mathbf{Q} -algebra, generated by the set $\{p^{1/2} \cdot \tau_p : p \text{ prime}\}$, subject only to the relations $(p^{1/2} \cdot \tau_p)^2 = p^2$. All \mathbf{Q} -algebra homomorphisms from this algebra to \mathbf{Q} are now given by $p^{1/2} \cdot \tau_p \mapsto \epsilon(p)p$, with ϵ ranging over $\{1, -1\}^{\mathcal{P}}$. Since we have $\{1, -1\}^{\mathcal{P}} \cong E_{\mathbf{Q}}$, this confirms that there are “just as many” such \mathbf{Q} -algebra homomorphisms as there are elements in $E_{\mathbf{Q}}$. The kernel of the homomorphism corresponding to ϵ is generated by the set $\{p^{1/2} \cdot \tau_p - \epsilon(p)p : p \text{ prime}\}$, in agreement with the proposition from Section 6 and the theory set forth above.

Explicitly writing down $K\langle\sqrt{K^*}\rangle^D$ for general K is a subject that we leave for further study.

15. Small subgroups of the group of radicals

Let K be a field of characteristic zero. Any computation with radicals involves only finitely many radicals, and these are all contained in a subgroup $A \subset \sqrt{K^*}$ that contains K^* as a subgroup of finite index. When describing and analysing algorithms involving radicals, we shall find it convenient to fix such a subgroup A at the outset, and restrict the computations to radicals that belong to A . It is then also convenient to have a version of the preceding theory in which the role of $\sqrt{K^*}$ is played by A .

As in Section 4, let first B be any multiplicatively written abelian group with the property that every finite subgroup of B is cyclic, and denote by \sqrt{B} a group of radicals of B . Let further A be an abelian group that contains B as a subgroup. Then there exists a B -embedding $A \rightarrow \sqrt{B}$ —that is, an injective group homomorphism that is the identity on B —if and only if the group A/B is torsion and every finite subgroup of A is cyclic. Suppose next that the group A/B is *finite*. Writing it as a direct sum of finitely many finite cyclic groups, we see that there exist a non-negative integer t , positive integers n_1, \dots, n_t , and elements $b_1, \dots, b_t \in B$, such that A may, as an abelian group containing B , be defined by t generators x_1, \dots, x_t and t relations $x_i^{n_i} = b_i$ ($1 \leq i \leq t$); necessarily, one has $\prod_{i=1}^t n_i = \#A/B$.

We apply this to $B = K^*$, writing “embedding” for “ K^* -embedding”. Let A be an abelian group that contains K^* as a subgroup of finite index. Then we define the K -algebra $K\{A\}$ to be the group ring $K[A]$ modulo the ideal generated by $\{a \cdot 1 - 1 \cdot a : a \in K^*\}$. If $t, n_1,$

\dots, n_t , and $b_1, \dots, b_t \in K^*$ are as above, then one has $K\{A\} \cong K[X_1, \dots, X_t]/(X_1^{n_1} - b_1, \dots, X_t^{n_t} - b_t)$ as K -algebras. Also, $K\{A\}$ has dimension $\#A/K^*$ as a vector space over K , a basis being given by a set of coset representatives for A modulo K^* .

Just as for $\sqrt{K^*}$, we can define the K -algebra $K\langle A \rangle$ in one of two equivalent ways: either as $K\{A\}$ modulo the ideal generated by $\{\sum_{\eta \in C} \eta : C \subset A \text{ is a subgroup of prime order with } C \not\subset K^*\}$, or as $S^{-1}K\{A\}$, where $S \subset K\{A\}$ is the multiplicative set generated by $\{a - 1 : a \in A, a \neq 1\}$. It turns out that $K\langle A \rangle$ is the zero ring unless (and only unless) every finite subgroup of A is cyclic. Assume now that the latter condition is satisfied. Then there is an embedding of A into $\sqrt{K^*}$, and given such an embedding one may identify $K\{A\}$ and $K\langle A \rangle$ with the sub- K -algebras of $K\{\sqrt{K^*}\}$ and $K\langle \sqrt{K^*} \rangle$ (respectively) that are generated by A . One can show that $K\langle A \rangle$ consists of all elements of $K\langle \sqrt{K^*} \rangle$ that are invariant under the subgroup $\text{Aut}_{A\sqrt{K^*}}$ of $\text{Aut}_{K^*\sqrt{K^*}}$. It is itself a Galois algebra over K with group $\text{Aut}_{K^*} A$ if and only if one has $\#\text{Aut}_{K^*} A = \dim_K K\langle A \rangle$, and if and only if there exists a positive integer n with $\#A[n] = n$ and $A^n \subset K^*$ (cf. Section 8). If A satisfies these equivalent conditions, we shall say that A is *Galois* over K^* .

If A is an abelian group that contains K^* as a subgroup of finite index, and every finite subgroup of A is cyclic, then one has $\dim_K K\langle A \rangle = \#A/K^* \cdot \prod_p (1 - 1/p)$, the product ranging over the set of prime numbers p for which $\#A[p] = p$ and $\#K^*[p] = 1$; this confirms that $K\langle A \rangle$ is not the zero ring.

As an example we consider the case in which the finite group A/K^* is *cyclic*, say of order n . Then we can write $A = K^* \cdot \langle a \rangle$ for some $a \in A$ with $a^n = b \in K^*$, and we have $K\{A\} \cong K[X]/(X^n - b)$. Denote by d the largest squarefree divisor of n for which $b \in K^{*d}$, say $b = c^d$ with $c \in K^*$. Then the condition that each finite subgroup of A be cyclic, is equivalent to c being *unique*, that is, to $\#K^*[d] = 1$. Let that be assumed. Then one has $K\langle A \rangle \cong K[X]/(f)$, where f equals $X^n - b$ divided by the least common multiple of the polynomials $X^{n/p} - c^{d/p}$, with p ranging over the prime numbers dividing d . Explicitly, one has $f = c^{\deg \Phi_d} \cdot \Phi_d(X^{n/d}/c)$, where Φ_d denotes the d th cyclotomic polynomial (see [18, Section 6.3]). One has $K\{A\} = K\langle A \rangle$ if and only if $f = X^n - b$, and if and only if $d = 1$.

16. Finite étale algebras

Let K be a field, and let F be a K -algebra. We say that F is *finite* if the vector space dimension $\dim_K F$ is finite, and if F is finite we say that F is *étale* if each element of F is a zero of a separable polynomial in $K[X]$; it suffices to require this for all elements in a set of K -algebra generators for F . For example, for $f \in K[X]$ the K -algebra $K[X]/(f)$ is finite étale if and only if f is separable. The algebras $K\{A\}$ and $K\langle A \rangle$ considered in the previous section are both finite étale. For each non-negative integer n , the K -algebra $K \times K \times \dots \times K$, with n factors K and componentwise operations, is finite étale; such an algebra is said to be *totally split*. For example, for a monic polynomial $f \in K[X]$ the K -algebra $K[X]/(f)$ is totally split if and only if there is a finite subset $S \subset K$ with $f = \prod_{a \in S} (X - a)$.

Each finite étale K -algebra F can be written as the product of finitely many finite separable field extensions of K . More specifically, if F is a finite K -algebra, then $\text{Spec } F$ is finite, the natural K -algebra homomorphism $F \rightarrow \prod_{\mathfrak{m} \in \text{Spec } F} F/\mathfrak{m}$ is *surjective*, and F is étale if and only if the latter map is an *isomorphism* and each of the fields F/\mathfrak{m} is separable over K . Actually *writing* a given finite étale K -algebra in this manner is not always easy; for $K[X]/(f)$ it amounts to factoring f into irreducible factors over K .

Let F be a finite étale K -algebra. Among all sub- K -algebras of F that are totally split, there is a unique maximal one, which we denote by F_{spl} . Under the isomorphism $F \rightarrow \prod_{\mathfrak{m} \in \text{Spec } F} F/\mathfrak{m}$, the subalgebra F_{spl} maps isomorphically to $\prod_{\mathfrak{m} \in \text{Spec } F} K$, so that one has $F_{\text{spl}} = \bigcap_{\mathfrak{m} \in \text{Spec } F} (K + \mathfrak{m})$.

If K is a finite field, then F_{spl} equals the *Berlekamp subalgebra* $\{a \in F : a^{\#K} = a\}$ of F , which plays an important role in algorithms for factoring polynomials over finite fields (see [31, Section 14.8]).

There are exactly $\#\text{Spec } F$ distinct K -algebra homomorphisms $F_{\text{spl}} \rightarrow K$, one for each $\mathfrak{m} \in \text{Spec } F$, and the various fields F/\mathfrak{m} are obtained as $F/\mathfrak{m} = F \otimes_{F_{\text{spl}}} K$. Thus, knowing F_{spl} and its maps to K enables us to write F as a product of fields. The role that F_{spl} plays for finite étale F is comparable to the role that the decomposition algebra plays in the case of Galois algebras. In fact, if F is a Galois algebra over K with a finite group G , then F is finite étale with $\dim_K F = \#G$, and the decomposition algebra F^D from Section 14 equals F_{spl} if and only if D is normal in G .

17. Fields generated by finitely many radicals

Let K be a field of characteristic zero, let Ω be an algebraically closed field containing K , and let A be an abelian group that contains K^* as a subgroup of finite index. We assume that every finite subgroup of A is cyclic. In Section 15 we defined the K -algebras $K\{A\}$ and $K\langle A \rangle$. In the present section we are interested in *field* extensions of K that may be said to be generated by A .

If $s: A \rightarrow \Omega^*$ is any embedding, then $K(sA)$ is a field extension of K that contains a subgroup isomorphic to A . If A is Galois over K^* , as defined in Section 15, then the image sA of A and the subfield $K(sA)$ of Ω are independent of the choice of s , so that with due care—as explained in Section 4—one can write $K(A)$ for $K(sA)$. However, for general A the notation $K(A)$ for $K(sA)$ is to be avoided, as it may refer to a field that is not even well-defined up to isomorphism. As an example, consider $K = \mathbf{Q}$, with $A = \mathbf{Q}^* \cdot \langle a \rangle$ containing \mathbf{Q}^* as a subgroup of index 16, and $a^{16} = -4$. Then there are embeddings $s, t: A \rightarrow \mathbf{C}^*$ for which the fields $\mathbf{Q}(sA)$ and $\mathbf{Q}(tA)$ are non-isomorphic. In this example, which is due to B. de Smit, both fields do have the same degree over \mathbf{Q} , namely 8. Below we shall see that, remarkably enough, this equality of degrees is a general phenomenon.

Returning to general K and A , we wish to understand the several fields $K(sA)$ through the ring $K\langle A \rangle$, and we shall do so through the spectrum $\text{Spec } K\langle A \rangle$ of that ring. For each embedding $s: A \rightarrow \Omega^*$, the kernel \mathfrak{m}_s of the induced K -algebra homomorphism $K\langle A \rangle \rightarrow \Omega$ belongs to $\text{Spec } K\langle A \rangle$, and one has $K(sA) \cong K\langle A \rangle / \mathfrak{m}_s$ (as field extensions of K). Also, each element of $\text{Spec } K\langle A \rangle$ is of the form \mathfrak{m}_s , and for two embeddings s, t one has $\mathfrak{m}_s = \mathfrak{m}_t$ if and only if there exists a K -automorphism τ of the field Ω with $t = \tau \circ s$. The upshot is that it suffices to study the set $\text{Spec } K\langle A \rangle$ as well as the K -algebras $K\langle A \rangle / \mathfrak{m}$, for $\mathfrak{m} \in \text{Spec } K\langle A \rangle$; since $K\langle A \rangle$ is finite étale over K , we may, by the previous section, equivalently study the maximal totally split subalgebra $K\langle A \rangle_{\text{spl}}$ and its maps to K .

The subalgebra $K\langle A \rangle_{\text{spl}}$ has a very elegant and useful Galois-theoretic description. Namely, think of A as being embedded into $\sqrt{K^*}$, so that we can view $K\langle A \rangle$ as a subring of $F = K\langle \sqrt{K^*} \rangle$. Here F is a Galois algebra over K with group $G = \text{Aut}_{K^*} \sqrt{K^*}$, and with a normal decomposition group $D = \text{Gal}(K\langle \sqrt{K^*} \rangle / K)$. As we mentioned in Section 15, the subalgebra $K\langle A \rangle$ consists of all elements of F that are invariant under the subgroup $\text{Aut}_A \sqrt{K^*}$ of G . Thanks to the fact that D is normal in G , one can now prove that inside F one has the equality

$$K\langle A \rangle_{\text{spl}} = F^D \cap K\langle A \rangle,$$

independently of the choice of the embedding $A \rightarrow \sqrt{K^*}$. Thus, $K\langle A \rangle_{\text{spl}}$ may be identified with the subring of elements of F that are invariant under the subgroup $D \cdot \text{Aut}_A \sqrt{K^*}$ of G . Also, from the fact that G/D is *abelian*, one deduces that the latter subgroup is *normal*, so that $K\langle A \rangle_{\text{spl}}$ is a totally split Galois algebra over K with a *finite abelian* group $G/(D \cdot \text{Aut}_A \sqrt{K^*})$. Both $\text{Spec } K\langle A \rangle_{\text{spl}}$ and $\text{Spec } K\langle A \rangle$ are therefore faithfully and transitively acted upon by the group $G/(D \cdot \text{Aut}_A \sqrt{K^*})$, which by the theorem of Section 13 is a quotient of the entanglement group E_K .

It is particularly striking that this beautiful state of affairs is valid *without* the condition that A be Galois over K^* . It is also pleasing that we did not just use the normality of D in G , but also the fact that G/D is abelian. As a consequence, one reads off that for each $\mathfrak{m} \in \text{Spec } K\langle A \rangle$, the field extension degree $[K\langle A \rangle/\mathfrak{m} : K]$ equals the group index $(D \cdot \text{Aut}_A \sqrt{K^*} : \text{Aut}_A \sqrt{K^*})$, independently of \mathfrak{m} . Thus, for all embeddings $s: A \rightarrow \Omega^*$, the fields $K(sA)$ have the same degree over K .

There is little doubt that the theory can, at this point, be continued with further results that assist in the explicit determination of $K\langle A \rangle_{\text{spl}}$ as a subalgebra of $K\langle A \rangle$. To begin with, one may consider the case in which A/K^* is cyclic. Adopting the notation $A = K^* \cdot \langle a \rangle$, $n = \#A/K^*$, $b = a^n \in K^*$, d, c, Φ_d from the end of Section 15, we have $b = c^d$, $\#K^*[d] = 1$, and $K\langle A \rangle \cong K[X]/(\Phi_d(X^{n/d}/c))$. Can one explicitly write down the maximal totally split subalgebra? Can one directly prove that all irreducible factors of $\Phi_d(X^{n/d}/c)$ over K have the same degree, and give a formula for that degree?

18. Algorithmic issues

In the remainder of these lectures, we address the problem of exactly computing with radicals. More specifically, let K be a field, and let Ω be an algebraically closed field containing K . We define the sequence of subfields $K_0 \subset K_1 \subset K_2 \subset \dots$ of Ω by putting $K_0 = K$, and by letting K_{i+1} , for $i \geq 0$, be the field obtained by adjoining to K_i all $a \in \Omega^*$ for which there is a positive integer n with $a^n \in K_i^*$; in particular, K_1 equals $K(\sqrt{K^*})$. We put $K_\infty = \bigcup_{i \geq 0} K_i$, the field of *nested radicals* over K^* . It is also called the *solvable closure* of K in Ω , since it is the union of all Galois extensions of K inside Ω that have solvable Galois groups. The problem is now how one may *represent* the elements of K_∞ in terms of those of K ; how, given representations of two elements of K_∞ , one may find representations for their *sum*, their *difference*, and their *product*; how one can *test equality* between two given elements, in case representations are not unique; and how one may construct the *inverse* of an element that is found to be different from 0.

Our concern is with *exact* computations, as opposed to *approximate* ones. Thus, one should not think of the field K as being given with a topology other than the discrete one.

The major challenge is to find a way of representing the elements of K_∞ that on the one hand does justice to those elements being expressible by means of “radicals”, and on the other hand makes the question whether two elements are equal into a well-defined one.

We are quite far from being able to offer complete solutions to the problems formulated. What we shall do, is report on algorithmic work that has been done and on problems to be faced. Our attitude will be practically oriented to the extent that we pay constant attention to the run times of the computational procedures under discussion, and that the concepts of computational complexity will guide us in drawing the line between the feasible and the infeasible.

We restrict our discussion to a special case that is still complicated enough to be interesting. Namely, we shall consider only the case in which K is an *algebraic number field*, that is, an extension field of finite degree of the field \mathbf{Q} of rational numbers. Number fields are of obvious interest, and they have a well-developed algorithmic theory that we can draw upon [7, 8, 20].

As is usual in the area of number-theoretic algorithms, we measure the run time of an algorithm on a certain input by the number of *bit operations* that it performs, and we are interested in bounding this run time as a function of the *length* of the input. Here the length of the input is simply the number of bits that it consists of. For example, the base field K itself will normally be part of the input; one may think of it as being numerically specified by means of an irreducible polynomial $f = \sum_{i=0}^n f_i X^i \in \mathbf{Z}[X]$ of positive degree n with the property that one obtains K from \mathbf{Q} by adjoining a zero a of f . The length of this portion of the input is then the number of bits required for spelling out the vector $(f_i)_{i=0}^n \in \mathbf{Z}^{n+1}$ of coefficients of f in full, including the zero entries, each entry being written in binary. If an element b of K also forms part of the input, then it is supposed to be numerically specified by means of the unique vector $(c_i)_{i=0}^{n-1} \in \mathbf{Q}^n$ for which $b = \sum_{i=0}^{n-1} c_i a^i$, each c_i being represented as a quotient of integers written in binary. Finite extensions of K and their elements are specified in a similar manner, the role of \mathbf{Q} now being played by K itself. For more information and background, one may consult [2, 20].

19. Computing with radicals

In the present section we discuss the possibility of designing an “efficient” algorithm for solving the following problem: given an algebraic number field K as well as an abelian group A containing K^* as a subgroup of finite index, decide whether there are *embeddings* s of A in the multiplicative group of an algebraically closed field Ω containing K , and provide explicit models for the fields $K(sA)$ obtained in this way. Here the field K and its elements are numerically specified as explained in the previous section, and A is numerically specified by a non-negative integer t , positive integers n_1, \dots, n_t , and elements b_1, \dots, b_t of K , as in Section 15; then A is, as an abelian group containing K^* , defined by t generators x_1, \dots, x_t , subject to the t relations $x_i^{n_i} = b_i$ ($1 \leq i \leq t$). We write n for the index $(A : K^*)$, which is given by $n = \prod_{i=1}^t n_i$. By an *embedding* $A \rightarrow \Omega^*$ we mean an injective group homomorphism that is the identity on K^* .

One may of course think of x_i as standing for $\sqrt[n_i]{b_i}$, but we avoid using the radical sign, not only because of the risk of improper usage but also because the same pair n_i, b_i may occur for different i .

There is a straightforward algorithm that solves the problem in time polynomial in the length on the input *and* the number n , and that uses very little of the theory developed so far. It proceeds by computing the ring $K\langle A \rangle$, as well as an explicit isomorphism of $K\langle A \rangle$ with a product of finitely many finite field extensions of K . To understand the algorithm, one should keep in mind that the ring $K\{A\}$, which is isomorphic to $K[X_1, \dots, X_t]/(X_1^{n_1} - b_1, \dots, X_t^{n_t} - b_t)$, is finite étale over K , and can therefore be written as the product of finitely many finite field extensions L of K ; one obtains the ring $K\langle A \rangle$ by discarding from this product those fields L for which the group homomorphism $A \rightarrow L^*$ induced by the projection map $K\{A\} \rightarrow L$ fails to be injective.

The algorithm proceeds by induction on t . For $t = 0$ one has $A = K^*$, and no work is necessary. For $t > 0$, let $A' \subset A$ be the subgroup of index n_t generated by x_1, \dots, x_{t-1} . Suppose that, inductively, the ring $K\langle A' \rangle$ has already been computed, and that it has been written as a product of finitely many finite field extensions L_i of K , with i ranging over a finite index set I . Write A'_i for the projection of A' to L_i ; this is a subgroup of L_i^* that contains K^* as a subgroup of index $n' = n/n_t$. Now one first factors the polynomial $X^{n_t} - b_t$ into irreducible factors in each of the rings $L_i[X]$, by means of standard algorithms (see [7, 31]). For each i , and for each of the irreducible factors $g \in L_i[X]$ that is found, one checks, by direct computation in the field $L_i[X]/(g)$, whether there is a positive integer m

with $m < n_t$ for which $(X \bmod g)^m$ belongs to one of the n' cosets of K^* in A'_i ; if such an integer m exists, then the pair i, g is discarded. The ring $K\langle A \rangle$ is now the product of all the fields $L_i[X]/(g)$ that remain. If no fields remain—which occurs, for example, if I is empty—then $K\langle A \rangle$ is the zero ring, there are no embeddings $A \rightarrow \Omega^*$, and one is finished. The fields that do remain are, as extensions of K , isomorphic to the fields $K(sA)$ that one is looking for.

The conclusion is that if n is small enough for the algorithm just outlined to be practical, then the theory that we developed can be largely bypassed. It is nevertheless worth investigating whether our theory can be used to improve the performance of the algorithm or to describe the fields $K(sA)$ in a uniform manner.

20. Polynomial-time algorithms

We retain the notation from the previous section. The algorithm that we described is not polynomial-time, as n enters only logarithmically into the length of the input. This makes itself felt when the index $n = (A : K^*)$ is very large: not so large that n cannot be written down, but too large to allow for algorithms to have run time at least n . Such a situation may occur if one adjoins a single radical x_1 that has a high power equal to an element b_1 of K^* , and also if all n_i equal 2, in which case the number $n = 2^t$ grows exponentially with the number t of square roots x_i that are adjoined.

It is, for very large n , not clear what one would mean by constructing an “explicit model” for one of the fields $K(sA)$ if that model is to be used for doing explicit computations in the field. The dimension of $K\{A\}$ over K equals n , and in most interesting cases the dimensions of $K\langle A \rangle$ and $K(sA)$ are not much smaller. Hence, for very large n , one is not even able to express “generic elements” of any of those rings on a K -basis. This severely limits the type of computations one is able to do. For example, expressing $1/(\sqrt[n]{2} - 1)$ on the \mathbf{Q} -basis $((\sqrt[n]{2})^i)_{i=0}^{n-1}$ of the field $\mathbf{Q}(\sqrt[n]{2})$ requires n non-zero coefficients. Likewise, when p_i denotes the i th prime number, then the inverse of $\sum_{i=1}^t \sqrt{p_i}$, when written down on the \mathbf{Q} -basis $(\prod_{i \in I} \sqrt{p_i})_{I \subset \{1, 2, \dots, t\}}$ of $\mathbf{Q}(\sqrt{p_1}, \dots, \sqrt{p_t})$, requires for each $t \geq 2$ at least $n/t = 2^t/t$ non-zero coefficients. These examples show that, if one insists on polynomial-time algorithms and on using a “natural” basis consisting of radicals, one should not ask for the computation of inverses of sums of radicals.

Nevertheless, there are in the present context a number of algorithmic problems for which one may study the possibility of finding polynomial-time algorithms. Here are a few. Given K, A , and a finite number of elements of A , decide whether these elements add

up to a unit of the ring $K\{A\}$; whether they add up to zero in the ring $K\langle A \rangle$; whether they add up to a unit in the ring $K\langle A \rangle$. Given K and A , decide whether the ring $K\langle A \rangle$ is non-zero, or, equivalently, whether there exists an embedding $s: A \rightarrow \Omega^*$; and if such an embedding exists, compute the common degree of all fields $K(sA)$ over K , with s ranging over all embeddings.

It is very unlikely that for the last problem a polynomial-time algorithm can currently be found. Namely, take $K = \mathbf{Q}$, restrict to odd n , and let A be the direct product of \mathbf{Q}^* and a cyclic group of order n (so $t = 1$, $n_1 = n$, $b_1 = 1$); then $[K(sA) : K]$ equals the value $\varphi(n)$ of the Euler phi-function in n . It is well-known that with knowledge of $\varphi(n)$, one can factor n into prime numbers by means of a probabilistic algorithm that runs in expected polynomial time (see [23]). Hence, any polynomial-time algorithm for the problem of computing $[K(sA) : K]$ would give rise to a fast probabilistic algorithm for factoring integers. Such an algorithm is not known to exist, and believed by many not to exist. Thus, in our last problem it may be reasonable to require that n is given in factored form, or to formulate the problem “modulo” the degree of a suitably defined “cyclotomic piece” of $K(sA)$.

It is plausible that at least for fixed K one can decide in polynomial time whether $K\langle A \rangle$ is the zero ring. In the case $K = \mathbf{Q}$, known algorithms related to the multiplicative group \mathbf{Q}^* are useful here (see [4, 5]). Corresponding “multiplicative” algorithms for number fields can be found in [3; 7, Section 6.5; 12; 13; 20, Section 5; 27]. Not all of these run in polynomial time for varying K .

The problem of describing $K\langle A \rangle_{\text{spl}}$ for given K and A leads to the non-algorithmic question of obtaining a “good” upper bound for $\dim_K K\langle A \rangle_{\text{spl}}$ as a function of t and K . One may also wish to compute E_K for given K ; the fact that for every algebraic number field K the group E_K is infinite, appears to be only a minor obstacle.

21. Nested radicals

Let K be an algebraic number field, let Ω be an algebraically closed field containing K , and let the sequence of subfields $K = K_0 \subset K_1 = K(\sqrt{K^*}) \subset K_2 \subset \dots$ of Ω be as defined in Section 18. In Sections 19 and 20 we discussed computational issues in the field K_1 , or rather in subfields of K_1 that are generated by finitely many radicals over K . In the present section and in Section 23 we discuss similar issues for the fields K_2, K_3, \dots

The obvious approach is to replace, in anything we said in Sections 19 and 20, the base field $K = K_0$ by the fields K_1, K_2, \dots in succession. There are, however, two essential

respects in which the fields K_1, K_2, \dots differ from $K = K_0$.

In the first place, the fields K_1, K_2, \dots contain the group μ of all roots of unity in Ω . By Kummer theory (see Section 5), this implies that K_{i+1} is, as a K_i -algebra, isomorphic to the ring $K_i\{\sqrt{K_i^*}\}$, and that the entanglement group E_{K_i} is trivial, for each $i \geq 1$. Thus, from a structural point of view, the first step $K_0 \subset K_1$ in the sequence is more complicated than any of the later steps $K_1 \subset K_2, K_2 \subset K_3, \dots$.

Secondly, none of the fields K_1, K_2, \dots is an algebraic number field: each of them has infinite degree over \mathbf{Q} . This gives rise to several complications. Algorithms for computing in an infinite algebraic extension L of \mathbf{Q} proceed most conveniently by restricting any given computation to a subfield of L that is of finite degree over \mathbf{Q} ; this subfield is chosen so as to contain the elements of L that form part of the input to the algorithm, and it is enlarged only as the need arises. That is the way we proceeded for the extension $L = K_1$ of $K = K_0$. For the other extensions $K_1 \subset K_2, K_2 \subset K_3, \dots$, it now becomes desirable to “approximate” *both* the extension field *and* the base field by finite degree subfields, which is a significant bother.

Infinite algebraic extensions of \mathbf{Q} offer additional algorithmic challenges that are relevant in the present context. One is to extend algorithms related to the multiplicative group (see the references in the previous section) to the infinite degree case. Another one, to be considered in the next section, is to find, for infinite algebraic extensions L of \mathbf{Q} , a good algorithm for factoring polynomials in $L[X]$ into irreducible factors in $L[X]$; in the finite case, such an algorithm played a crucial role in Section 19.

There are uncountably many pairwise non-isomorphic infinite algebraic field extensions of \mathbf{Q} , so they cannot all be numerically specified by a finite number of bits. Thus, it is not possible to pose algorithmic questions for all of these fields in a meaningful way, and it becomes necessary to restrict to fields that are described in some special manner, such as the fields K_i considered above, and their union $K_\infty = \bigcup_{i \geq 0} K_i$. Depending on the way in which the field L is specified, it may not even be obvious whether an algorithm for factoring in $L[X]$ exists at all, let alone a “good” one.

22. Algorithms for large number fields

Let K be an algebraic number field, let $K \subset L$ be a possibly infinite Galois extension of K , and let Ω be an algebraically closed field containing L . If M is any algebraic extension of K , then there is a field homomorphism $s: M \rightarrow \Omega$ that is the identity on K , and thanks to the fact that L is Galois over K , the subfield $s^{-1}L$ of M is independent of s ; we shall simply write $L \cap M$ for this field.

Proposition. *Let K and L be as above. Suppose that there exists an algorithm for solving one of the following three computational problems:*

- (i) *given a finite field extension M of K , and an element $a \in M$, decide whether $a \in L \cap M$;*
- (ii) *given a finite field extension M of K , determine $L \cap M$;*
- (iii) *given an irreducible polynomial $f \in K[X]$, find an irreducible factor of f in $L[X]$.*

Then there are also algorithms for solving the other two.

In this proposition, one should think of K as specified by a defining polynomial over \mathbf{Q} (see Section 18). The field M is, as part of the input to problems (i) and (ii), specified by a defining polynomial over K . In problem (ii), one asks for a K -basis of $L \cap M$ or, equivalently, for an element $b \in M$ with $K(b) = L \cap M$. In problem (iii), one asks for a finite extension $K \subset M$ as well as a monic factor g of f in $M[X]$ such that for any s as above, the induced ring homomorphism $M[X] \rightarrow \Omega[X]$ (mapping X to X) sends g to an irreducible factor of f in $L[X]$. (As s varies, *all* monic irreducible factors of f in $L[X]$ will be obtained in this way from g , so asking for a single factor suffices.) Note that the proposition talks only about the *existence* of algorithms, not about their efficiency.

The proof of the proposition is an exercise in field theory. It can be outlined as follows. Suppose first that one can do (ii). Then one can trivially do (i), and one can do (iii) by applying (ii) to $M = K(a)$ with $f(a) = 0$ and determining the irreducible polynomial g of a over $L \cap M$; it is one of the irreducible factors of f in $L[X]$. Next suppose that one can do (i). Then one can do (ii) by choosing $a \in M$ with $M = K(a)$, factoring the irreducible polynomial f of a over K into irreducible factors in $M[X]$, and testing, for each of the monic factors of f in $M[X]$ that are divisible by $X - a$, whether its coefficients are in $L \cap M$; the coefficients of the smallest degree factor with this property generate $L \cap M$ as a field extension of K . Finally, suppose that one can do (iii). Then one can do (ii) by choosing $a \in M$ such that $M = K(a)$ and using (iii) to find a monic irreducible factor g in $L[X]$ of the irreducible polynomial of a over K ; if E is the field generated by the coefficients of g ,

then the K -algebra homomorphism $\phi: M = K(a) \rightarrow E[X]/(g)$ sending a to $(X \bmod g)$ is actually an *isomorphism*, and one has $L \cap M = \phi^{-1}E$. This completes our outline of the proof of the proposition.

We note that, if there is an algorithm for solving one of (i), (ii), and (iii), and hence for the others, then we can also factor in $L[X]$. That is, we can meaningfully write down and recognize polynomials with coefficients in L , and there is an algorithm for factoring them into irreducible factors in $L[X]$. To do this, one forms the product of the K -conjugates of the given polynomial, and one applies (iii) to its irreducible factors in $K[X]$.

A version of the proposition applying to *polynomial-time* algorithms would have great interest. Three of the four reductions that we gave in the proof are polynomial-time; the remaining one, reducing (ii) to (i), is not, and one may wonder whether one can design a polynomial-time reduction of (ii) to (i). It would suffice to have a polynomial-time algorithm that, given $K \subset M$, lists all minimal elements of the set $\{E : E \text{ is a subfield of } M \text{ containing } K, E \neq K, \text{ and there is a Galois extension } L \text{ of } K \text{ with } E = L \cap M\}$. This leads to the following problem on finite groups.

Problem. *Decide whether there is a positive real number c with the following property. Let G be a finite group, let $H \subset G$ be a subgroup with $H \neq G$, and let \mathcal{I} be the set of subgroups $I \subset G$ with $I \neq G$ for which there exists a normal subgroup $N \subset G$ with $I = N \cdot H$. Then the number of elements of \mathcal{I} that are maximal under inclusion is at most $(G : H)^c$.*

23. Abelian and solvable fields

For certain special classes of Galois extensions of number fields there do exist polynomial-time algorithms for solving the three problems listed in the proposition from the previous section.

For an algebraic number field K contained in an algebraically closed field Ω , let $K(\mu)$ be the maximal cyclotomic extension of K inside Ω (see Section 9), let K^{ab} be the maximal Galois extension of K inside Ω that has an abelian Galois group, and let K_∞ , as in Section 18, be the solvable closure of K in Ω .

Proposition. *There are polynomial-time algorithms that, given an algebraic number field K and a finite field extension M of K , compute the subfields $K(\mu) \cap M$, $K^{\text{ab}} \cap M$, and $K_\infty \cap M$ of M .*

An outline of the proof is as follows. In order to compute $K^{\text{ab}} \cap M$ for given M and K , one starts by computing the intersection of all K -conjugates of M ; this can be done in polynomial time, and it gives rise to the largest Galois extension E of K that is contained in M . Next one computes the group $G = \text{Gal}(E/K)$ and its commutator subgroup $[G, G]$ in a straightforward way. Then the field $K^{\text{ab}} \cap M$ that one is looking for equals the field $E^{[G, G]}$ of invariants of $[G, G]$. To compute $K(\mu) \cap M$ one can now use the formula $K(\mu) \cap M = K \cdot (\mathbf{Q}^{\text{ab}} \cap M)$; it is a consequence of the Kronecker-Weber theorem, which asserts that \mathbf{Q}^{ab} equals $\mathbf{Q}(\mu)$.

The algorithm for computing $K_\infty \cap M$ is a bit more involved. First one writes the finite étale K -algebra $M \otimes_K M \otimes_K M \otimes_K M$ as a product of finitely many field extensions E of K . Next one determines, for each E , by a straightforward method similar to the determination of $K^{\text{ab}} \cap M$ above, the largest intermediate field E' of $K \subset E$ that is a Galois extension of K with a solvable Galois group, and one also determines the subfield $M_E = \{a \in M : \text{the image of } a \otimes 1 \otimes 1 \otimes 1 \text{ in } E \text{ lies in } E'\}$ of M . If all fields M_E equal K , then one has $K_\infty \cap M = K$. Otherwise, one lets F be the composite of all fields M_E , and computes $K_\infty \cap M$ by the recursive formula $K_\infty \cap M = F_\infty \cap M$. The correctness proof of this algorithm depends on the following result from group theory (see [28]): *if G is a solvable group that acts primitively on a finite set X , then there is a subset $Y \subset X$ with $\#Y \leq 4$ such that each $\sigma \in G$ that acts as the identity on Y , acts as the identity on all of X .* This concludes the outline of the proof of the proposition.

The algorithm for $K_\infty \cap M$ improves upon an algorithm of Landau and Miller [17], which decides in polynomial time, for given K , M , and $a \in M$, whether $a \in K_\infty \cap M$.

For K and Ω as above, let the subfields K_1, K_2, \dots of Ω be as defined in Section 18. Cotner [9] showed that there is an algorithm that given K , M , and $a \in K_\infty \cap M$, determines the least non-negative integer i with $a \in K_i \cap M$; this is the *nesting depth* of a as a radical expression over K . The main ingredient of Cotner's algorithm is a method for computing $K_1 \cap M$. This method depends on investigating the ramification in the extension $K \subset K_1$, and it does not run in polynomial time. One may wonder whether the method used in the proof of the above proposition may be extended to a polynomial-time method for computing $K_1 \cap M$, or indeed for computing all of the fields $K_i \cap M$.

Acknowledgments

The author is grateful to I. M. Isaacs and T. R. Wolf for pointing out the work of Seress, to P. Stevenhagen for organizing the special session accompanying the lecture series, to B. de Smit and W. J. Palenstijn for permission to quote from their unpublished work in Section 13, to B. de Smit and W. Bosma for many helpful discussions, to W. J. Palenstijn for typographical assistance and for creating the illustration on the title page, and to the *Nieuw Archief voor Wiskunde* for permission to reproduce it.

References

1. T. Albu, *Cogalois theory*, Marcel Dekker, Inc., New York, 2003.
2. E. Bach, J. Shallit, *Algorithmic number theory*, volume 1, MIT Press, Cambridge, Mass., 1996.
3. D. J. Bernstein, *Fast ideal arithmetic via lazy localization*, Algorithmic number theory (Talence, 1996), 27–34, Lecture Notes in Comput. Sci. **1122**, Springer, Berlin, 1996.
4. D. J. Bernstein, *Factoring into coprimes in essentially linear time*, J. Algorithms **54** (2005), 1–30.
5. D. J. Bernstein, H. W. Lenstra, Jr., J. Pila, *Detecting perfect powers by factoring into coprimes*, Math. Comp., to appear.
6. J. W. S. Cassels, A. Fröhlich (eds), *Algebraic number theory*, Academic Press, London, 1967.
7. H. Cohen, *A course in computational algebraic number theory*, Springer, Berlin, 1993.
8. H. Cohen, *Advanced topics in computational number theory*, Springer, New York, 2000.
9. C. F. Cotner, *The nesting depth of radical expressions*, Ph. D. thesis, University of California at Berkeley, 1995.
10. B. de Smit, W. J. Palenstijn, in preparation.
11. L. Fuchs, *Infinite abelian groups*, vol. I, Academic Press, New York, 1970.
12. G. Ge, *Algorithms related to multiplicative representations of algebraic numbers*, Ph. D. thesis, University of California at Berkeley, 1993.
13. G. Ge, *Recognizing units in number fields*, Math. Comp. **63** (1994), 377–387.
14. M. Honsbeek, *Radical extensions and Galois groups*, Ph. D. thesis, Radboud Universiteit Nijmegen, 2005.
15. C. Hooley, *On Artin’s conjecture*, J. Reine Angew. Math. **225** (1967), 209–220.
16. M. Kneser, *Lineare Abhängigkeit von Wurzeln*, Acta Arith. **26** (1975), 307–308.

17. S. Landau, G.L. Miller, *Solvability by radicals is in polynomial time*, J. Comput. System Sci. **30** (1985), 179–208.
18. S. Lang, *Algebra*, revised third edition, Springer, New York, 2002.
19. D. H. Lehmer, E. Lehmer, *Heuristics, anyone?*, Studies in Mathematical Analysis and Related Topics, Stanford University Press, 1962, 202–210; *Selected papers of D. H. Lehmer*, The Charles Babbage Research Center, St. Pierre, 1981, vol. I, 357–365.
20. H. W. Lenstra, Jr., *Algorithms in algebraic number theory*, Bull. Amer. Math. Soc. **26** (1992), 211–244.
21. H. W. Lenstra jr., *Profinite Fibonacci numbers*, Nieuw Arch. Wisk. (5) **6** (2005), 297–300.
22. H. W. Lenstra, Jr., P. Stevenhagen, *Chebotarëv and his density theorem*, Math. Intelligencer **18** (2) (1996), 26–37.
23. D. Long, *Random equivalence of factorization and computation of orders*, Princeton U. Dept. Elec. Eng. and Comp. Sci. Technical Report **284** (1981).
24. W. J. Palenstijn, *Galois action on division points*, Master’s thesis, Universiteit Leiden, 2004 (<http://www.math.leidenuniv.nl/scripties/palenstijn.pdf>).
25. L. Ribes, P. Zalesskii, *Profinite groups*, Springer, Berlin, 2000.
26. A. Schinzel, *Abelian binomials, power residues and exponential congruences*, Acta Arith. **32** (1977), 245–274.
27. R. J. Schoof, *Computing Arakelov class groups*, J. P. Buhler, P. Stevenhagen (eds), *Surveys in algorithmic number theory*, Mathematical Sciences Research Institute Publications, Cambridge University Press, to appear.
28. Á. Seress, *The minimal base size of primitive solvable permutation groups*, J. London Math. Soc. (2) **53** (1996), 243–255.
29. P. Stevenhagen, *The correction factor in Artin’s primitive root conjecture*, J. Théor. Nombres Bordeaux **15** (2003), 383–391.
30. J. Tate, *Les conjectures de Stark sur les fonctions L d’Artin en $s = 0$* , Birkhäuser, Boston, 1984.
31. J. von zur Gathen, J. Gerhard, *Modern computer algebra*, Cambridge University Press, 1999.
32. J. S. Wilson, *Profinite groups*, Clarendon Press, Oxford, 1998.
33. J. W. Wrench, Jr., *Evaluation of Artin’s constant and the twin-prime constant*, Math. Comp. **15** (1961), 396–398.