Kandidátusi Értekezés
(Ph. D. Thesis submitted to the Hungarian Academy of Sciences)

Algorithms for algebras over global fields

Ivanyos Gábor

1996

Abstract

The main results in this dissertation concern the computational complexity of structural decomposition problems in finite dimensional associative algebras over global fields (algebraic number fields and global function fields, i.e., function fields of plane algebraic curves over finite fields).

Polynomial time algorithms for isolating the radical and finding the simple components of the semisimple part of an algebra over a global function field are presented.

We propose a method for computing the dimension of minimal one-sided ideals of a simple algebra over a global field. The method is based on computing a *maximal order* in the algebra, a non-commutative analogue of the ring of algebraic integers in a number field. The algorithm makes oracle calls to factor integers in the number fields case.

A generalization of the LLL basis reduction algorithm is used to demonstrate that computing a maximal order in an algebra isomorphic to $\mathrm{M}_2(\mathbb{Q})$ is equivalent to finding an explicit isomorphism with $\mathrm{M}_2(\mathbb{Q})$.

We also present some applications, such as an efficient membership test in commutative matrix groups as well as a polynomial time method for computing dimensions of irreducible representations of finite groups over number fields.

Some results are also valid in more general contexts. For example, a deterministic polynomial time method for finding a maximal toral subalgebra in a semisimple algebra is presented as an application of a method for computing a Cartan subalgebra in a Lie algebra.

Acknowledgements

I am grateful to my scientific collaborators, in particular to László Babai, Robert Beals, Jin-yi Cai, Arjeh M. Cohen, Willem A. de Graaf, Eugene M. Luks, Lajos Rónyai, Ágnes Szántó, and David B. Wales, who were coauthors of the papers this dissertation is based on. I am especially indebted to my supervisor Lajos Rónyai for his invaluable advice.

My thanks also go to my colleagues at the Computer and Automation Institute of the Hungarian Academy of Sciences, in particular to the collective led by János Demetrovics for ensuring support and the excellent working atmosphere making my research possible. I am grateful as well to the collective around András Recski, (the Department of Mathematics and Computer Science of the Faculty of Electrical Engineering and Informatics at the Technical University of Budapest) for hosting me as a corresponding research student.

Finally, but not least I would like to thank my family: my wife, my children and my parents for their sacrifice and the warm, loving atmosphere they provided.

Contents

Chapter 1

Introduction

There is a considerable interest in computations with finite dimensional algebras as they emerge naturally in several fields of mathematics and its applications. For example, understanding the structure of associative algebras is an important tool in the theory of matrix groups. Since decomposition of problems into smaller ones plays an important role in designing efficient algorithms, structural decomposition of algebras serve as a general tool in solving computational problems related to matrices. In this thesis we present some recent developments in the area of symbolic computations related to the structure of finite dimensional algebras.

The organization of the dissertation is as follows. In this introductory chapter we give a short summary of the mathematical background of problems addressed (Section 1.1) as well as a description of the computational model and the basic algorithmic ingredients of the methods presented later (Section 1.2). A short survey of the most important results obtained by other authors is given in Section 1.3. Throughout, the term algebra is reserved for a finite dimensional associative algebra over a field.

In Chapter 2, based on part of the paper [BBCIL], joint work with László Babai, Robert Beals, Jin-yi Cai, and Eugene M. Luks, we present an application (Theorem 2.1.1) of algebra decompositions over number fields to a strong membership test for commutative matrix groups.

Efficient algorithms are known over finite fields and algebraic number fields for computing the radical and for finding the simple components of the radical-free part of algebras. Here, in Chapters 3 and 4, we extend these results to algebras over global function fields, i.e., finite algebraic extensions of the field $\mathbb{F}_q(X)$ of rational functions over the field $\mathbb{F}_q$ consisting of q elements. It turns out, however, that the methods admit natural extensions to algebras over algebraic function fields over $\mathbb{F}_q$ (finite extensions of $\mathbb{F}_q(X_1, \ldots, X_m)$), therefore the results are presented in this more general setting.

In Chapter 3, based on parts of the papers [IRSz] (joint work with Lajos Rónyai and Ágnes Szántó) and [CIW] (joint work with Arjeh M. Cohen and David B. Wales), a polynomial time algorithm for computing the radical of algebras over global function fields is presented (Corollary 3.4.6 to Theorem 3.4.5). The method is based on Theorem 3.1.4, a characterization of the radical of algebras over arbitrary fields of positive characteristic. This extends a result of Rónyai [Ró2], who gave a characterization for finite ground fields. In contrast to the method in [Ró2], which was based on certain functions obtained from lifting matrices over the finite prime field $\mathbb{F}_p$ to matrices over $\mathbb{Z}$, we use certain coefficients of the characteristic polynomial. However, in the finite case, the functions in both methods turn out to be essentially the the same (Proposition 3.2.3), whence the methods presented here give alternatives to some details of Rónyai's original method. As demonstrated in Theorem 3.3.2, the collection of these coefficients generalize the role of the trace in representation theory of semisimple algebras over fields of characteristic zero.

In Chapter 4, based on part of the paper [IRSz], a deterministic polynomial time method which is allowed to make oracle calls to factor polynomials over the prime field (an f-algorithm) for computing the Wedderburn decomposition of semisimple algebras over global function fields is presented (Corollary 4.1.5 to Theorem 4.1.3). The method is an improved analogue of the algorithm of Gianni, Miller and Trager [GMT], which was designed for decomposition of algebras over number fields.

Chapter 5 is mostly about Lie algebras. Cartan subalgebras are extremely important in the classification of (simple) Lie algebras. Here, based on the paper [GIR], joint work with Willem A. de Graaf and Lajos Rónyai, we present deterministic polynomial time algorithms for finding Cartan subalgebras in Lie algebras over sufficiently large fields (Theorem 5.4.1) as well as in Lie algebras over finite fields belonging to an important subclass (Theorem 5.4.2). How this result can be applied to derandomize several randomized methods for associative algebras is shown in Section 5.5.

Chapter 6, based on the paper [IR], joint work with Lajos Rónyai, is devoted to results related to the structure of simple algebras over global fields. The methods are based on certain noncommutative generalizations of ideas from algebraic number theory. The central result, stated in Theorem 6.4.2, is a deterministic polynomial time method allowed to make oracle calls to find prime factors of integers and to factor polynomials over finite fields (an ff-algorithm), that finds a maximal order (a noncommutative analogue of the ring of algebraic integers in number fields) in a semisimple algebra over a number field. An interesting application (Theorem 6.5.5) is a polynomial time algorithm for computing the dimensions of the irreducible constituents of a representation of a finite group over a number field. Analogous f-algorithms for computing maximal orders and indices in algebras over global function fields are also discussed.

In Chapter 7, based on the paper [ISz], joint work with Ágnes Szántó, we address the problem of complexity of finding zero divisors in maximal orders in simple algebras. We present a polynomial time method for central simple algebras of dimension four over the field of rationals (Theorem 7.2.1). The method is based on a generalization of the celebrated basis reduction procedure by A. K. Lenstra, H. W. Lenstra and L. Lovász [LLL] to the case of indefinite quadratic forms, discussed in Section 7.1.

We conlude with some open problems in Chapter 8.

1.1 Basic facts and definitions

(Nonassociative) algebras

A linear space $\mathcal{A}$ over the field K is an *algebra* over K if it is equipped with a binary, K-bilinear operation $(x, y) \mapsto xy$ (called multiplication). $\mathcal{A}$ is *associative* if

$$x(yz) = (xy)z \text{ holds for every } x, y, z \in \mathcal{A}.$$

Throughout this thesis we reserve the term *algebra* for associative algebras. In order to distinguish from the general case, for not necessarily associative algebras we use the term *nonassociative algebra*. We restrict ourselves to *finite dimensional* K-algebras. Besides associative algebras, important examples are *Lie algebras*. In the Lie case, we use the traditional bracket notation for multiplication. A nonassociative K-algebra $\mathcal{L}$ with multiplication $(x, y) \mapsto [x, y]$ is a Lie algebra over K if

$$[x, y] = -[y, x] \quad \text{(anticommutativity)}$$

and

$$[[x, y], z] + [[y, z], x] + [[z, x], y] = 0 \quad \text{(the Jacobi identity)}$$

hold for every $x, y, z \in \mathcal{L}$.

We say that two elements $x, y \in \mathcal{A}$ *commute* if $xy = yx$. A nonassociative algebra $\mathcal{A}$ is called *commutative* (or *abelian*), if every pair of its elements commute. A K-subspace $\mathcal{B}$ of $\mathcal{A}$ is a subalgebra of $\mathcal{A}$ ($\mathcal{B} \leq \mathcal{A}$ in notation) if it is closed under multiplication. A K-subspace L of $\mathcal{A}$ is a *left ideal* of $\mathcal{A}$ if $yx \in L$ holds whenever $x \in L$ and $y \in \mathcal{A}$. A *right ideal* is defined in an analogous fashion. A K-subspace I of $\mathcal{A}$ is an *ideal* (two-sided ideal) of $\mathcal{A}$ if I is both left and right ideal of $\mathcal{A}$. We use the standard notation $I \lhd \mathcal{A}$. Note that in commutative algebras as well as in Lie algebras the notions of left ideal, right ideal and ideal coincide.

For nonassociative K-algebras $\mathcal{A}$ and $\mathcal{B}$ a K-linear map $\phi : \mathcal{A} \to \mathcal{B}$ is a *homomorphism* if it preserves multiplication. The kernel $\ker \phi$ is an ideal in $\mathcal{A}$, while the image $\operatorname{im} \phi$ is a subalgebra of $\mathcal{B}$. An *isomorphism* is a bijective homomorphism. If $I \lhd \mathcal{A}$ is an ideal, then the factor space $\bar{\mathcal{A}} = \mathcal{A}/I$ inherits the multiplication of $\mathcal{A}$ in the natural way: $(x+I)(y+I) \subseteq xy + I$ holds for every $x, y \in \mathcal{A}$. Here we used the standard notation for extending operations to complexes (subsets): if $X, Y \subseteq \mathcal{A}$ then $X * Y$ stands for the subset $\{x * y | x \in X, y \in Y\}$, where $*$ is one of the operations on $\mathcal{A}$. The map $\phi : x \mapsto x + I$ is a homomorphism $\mathcal{A} \to \bar{\mathcal{A}}$, called the *natural map*. We have $\ker \phi = I$ and $\operatorname{im} \phi = \bar{\mathcal{A}}$.

A nonassociative algebra $\mathcal{A}$ is *simple* if it has only trivial ideals (i.e., (0) and $\mathcal{A}$) and $\mathcal{A} \neq (0)$. We say that $\mathcal{A}$ is the *direct sum* of its (left) ideals $\mathcal{A}_1, \ldots, \mathcal{A}_r$ (written as $\mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_r$) if $\mathcal{A}$ is the direct sum of these linear subspaces.

Associative algebras

Throughout this subsection $\mathcal{A}$ is a finite dimensional associative algebra over the field K. Because of associativity, the product $x_1 x_2 \cdots x_r$ of r elements $x_1, x_2 \ldots, x_r \in \mathcal{A}$ can be defined in a straightforward way.

The centralizer $\mathrm{C}_{\mathcal{A}}(X)$ of a subset X of $\mathcal{A}$ is the set consisting of elements of $\mathcal{A}$ commuting with every element of the subset X. Obviously, $\mathrm{C}_{\mathcal{A}}(X) \leq \mathcal{A}$. The center $\mathrm{C}(\mathcal{A})$ of $\mathcal{A}$ is the subalgebra $\mathrm{C}_{\mathcal{A}}(\mathcal{A})$.

An element $e \in \mathcal{A}$ is called an *identity element* if

$$ex = xe = x \text{ holds for every } x \in \mathcal{A}.$$

If $\mathcal{A}$ admits an identity element then the identity element is known to be unique and denoted by $1_{\mathcal{A}}$ or simply by 1. Note that if $\mathcal{A}$ has no identity element then we can adjoin one using the *Dorroh extension*: Let $\mathcal{A}' = Ke \oplus \mathcal{A}$ as vector spaces with multiplication defined by

$$(\alpha e + x)(\beta e + y) = \alpha\beta e + \beta x + \alpha y + xy.$$

It is easy to see that $\mathcal{A}'$ is an associative K-algebra with identity element e such that $\mathcal{A}$ is an ideal in $\mathcal{A}'$. The left, right, or two-sided ideals of $\mathcal{A}'$ are $\mathcal{A}'$ and those of $\mathcal{A}$.

A pair of nonzero elements $x, y \in \mathcal{A}$ is a pair of *zero divisors* in $\mathcal{A}$ if $xy = 0$. From the assumption that $\mathcal{A}$ is finite dimensional it follows that $x \in \mathcal{A}$ is the left member of a pair of zero divisors if and only if x is the right member of a pair of zero divisors. We call such an x a zero divisor. It turns out, that x is a zero divisor iff for the left ideal $\mathcal{A}x$ we have $\mathcal{A} > \mathcal{A}x$ and iff for the right ideal $x\mathcal{A}$ we have $\mathcal{A} > x\mathcal{A}$. Algebras without zero divisors (called *division algebras* over K or *skewfield extensions* of K) are obviously simple and a commutative algebra $\mathcal{A}$ is simple if and only if $\mathcal{A}$ admits no zero divisors. Therefore every finite dimensional commutative simple algebra over K is isomorphic to a finite extension field of K.

If V is an n-dimensional vector space over K then $\mathrm{End}_K(V)$, the algebra of K-linear transformations of V with the usual operations, is a simple K-algebra of dimension n^2. By choosing a basis, we can identify V with the space K^n of column vectors of length n and $\mathrm{End}_K(V)$ with $\mathrm{M}_n(K)$, the algebra of n by n matrices over K with the usual matrix operations.

Subalgebras of $\mathrm{End}_K(V)$ or, equivalently, those of $\mathrm{M}_n(K)$, called *matrix algebras*, appear to be typical examples of associative algebras. If $\mathcal{A}$ has an identity element then $\mathcal{A}$ can be efficiently embedded as a subalgebra of $\mathrm{M}_n(K)$, where $n = \dim_K \mathcal{A}$. This is easily seen using the (left) *regular representation*. For $x \in \mathcal{A}$ we define the linear map $L_x : \mathcal{A} \to \mathcal{A}$, called the left action of x on $\mathcal{A}$ as $L_x(y) = xy$ for every $y \in \mathcal{A}$. It is straightforward that $x \mapsto L_x$ is an algebra homomorphism of $\mathcal{A}$ to the algebra of linear transformations of the linear space $\mathcal{A}$. Moreover, if $\mathcal{A}$ has an identity element then $x \mapsto L_x$ is an injective map. If $\mathcal{A}$ has no identity element then the regular representation of the Dorroh extension induces an embedding $\mathcal{A} \to \mathrm{M}_{n+1}(K)$.

Matrix algebras arise naturally in problems related to *common invariant subspaces* of matrices. Let $X \subseteq \mathrm{End}_K(V)$ be a set (e.g., a group) of linear transformations of the finite dimensional linear space V. Obviously, if W is an X-invariant subspace (i.e., $xW \subseteq W$ for every $x \in X$) then W is also $\mathcal{A}$-invariant where $\mathcal{A}$ is the subalgebra of $\mathrm{End}_K(V)$ generated by X and the identity (the smallest subalgebra containing $X \cup \{\mathrm{Id}_V\}$). The centralizer algebra $\mathrm{C}_{\mathrm{End}_K(V)}(X)$ plays an important role in problems related to direct decompositions. To be more specific, decompositions of $\mathrm{C}_{\mathrm{End}_K(V)}(X)$ into direct sums of left ideals correspond to decompositions of V into direct sums of X-invariant subspaces. Other important examples of finite dimensional algebras are *group algebras* of finite groups.

An element $x \in \mathcal{A}$ is *nilpotent* if $x^N = 0$ for some positive integer exponent N. For a positive integer j and a subset $X \subseteq \mathcal{A}$ we denote the set $\{x_1 \cdots x_j | x_1, \ldots, x_j \in X\}$ by X^j. It is straightforward to see that if X is a K-subspace (subalgebra, left ideal, right ideal, ideal) of $\mathcal{A}$ then X^j is a K-subspace (subalgebra, left ideal, right ideal, ideal, resp.) as well. A subalgebra $\mathcal{B}$ is called *nilpotent* if $\mathcal{B}^N = 0$ for some integer $N > 0$. This in turn is equivalent to that $\mathcal{B}^N = 0$ for some integer $0 < N \leq \dim_K \mathcal{B} + 1$. It is known that a subalgebra $\mathcal{B}$ is nilpotent if and only it consists of nilpotent elements. There exists a largest nilpotent ideal of $\mathcal{A}$, called the *radical* of $\mathcal{A}$ and denoted by $\mathrm{Rad}(\mathcal{A})$. There are several characterizations of the radical, such as the intersection of the maximal ideals, or the set of *strongly nilpotent* elements (where x is said to be strongly nilpotent if x, xy, yx

are nilpotent for every $y \in \mathcal{A}$), etc. Note that the two-sided characterizations above could be replaced by analogous left-sided or right-sided ones.

$\mathcal{A}$ is called *semisimple* if $\mathrm{Rad}(\mathcal{A}) = \{0\}$. It turns out that the factoralgebra $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ is semisimple. We call $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ the semisimple part (or radical-free part) of $\mathcal{A}$. There is a very strong and useful characterization of semisimple algebras, due to Wedderburn.

Wedderburn's Theorem. *Let $\mathcal{A}$ be a finite dimensional algebra over the field K.*
(i) $\mathcal{A}$ is semisimple if and only if $\mathcal{A}$ is a direct sum of simple algebras

$$\mathcal{A} = \mathcal{A}_1 \oplus \cdots \oplus \mathcal{A}_r,$$

where the $\mathcal{A}_i$ are the only minimal nontrivial ideals of $\mathcal{A}$.
(ii) $\mathcal{A}$ is simple if and only if

$$\mathcal{A} \cong \mathrm{M}_t(D),$$

where D is a division algebra over K and t is a positive integer.

Let $\mathcal{A}$ be semisimple. We keep ourselves to the notation of the theorem. The minimal ideals $\mathcal{A}_1, \ldots, \mathcal{A}_r$ are also called the *simple components* of $\mathcal{A}$, and the decomposition (i) in the theorem is the *Wedderburn decomposition* of $\mathcal{A}$. We remark that the Wedderburn decomposition of the center corresponds to the decomposition of $\mathcal{A}$: the minimal ideals of $\mathrm{C}(\mathcal{A})$ are $\mathrm{C}(\mathcal{A}_1), \ldots, \mathrm{C}(\mathcal{A}_r)$.

A semisimple algebra $\mathcal{A}$ necessarily admits an identity element. In that case we identify K with the subalgebra $K1_{\mathcal{A}} \leq \mathrm{C}(\mathcal{A})$. An algebra $\mathcal{A}$ is *central* over K if $\mathrm{C}(\mathcal{A}) = K$. Every simple algebra is central over its center, which is a finite extension field of K. Assume that $\mathcal{A}$ is central simple over K. We know that $\dim_K \mathcal{A}$ is a square, say n^2, the number t in Wedderburn's theorem (ii) is a divisor of n, while D is a central division algebra over K of dimension $(\frac{n}{t})^2$. The number $\frac{n}{t}$ is called the *index* of $\mathcal{A}$. The minimal left ideals of $\mathcal{A}$ have dimension $\frac{n^2}{t}$ over K.

The *minimal polynomial* of an element a of a K-algebra $\mathcal{A}$ with identity is the monic polynomial $f \in K[X]$ such that $f(a) = 0$ and f is of minimum degree among the polynomials satisfying this property. (For a polynomial $g(X) = \sum_{i=0}^{d} \alpha_i X^i$, $g(a)$ is defined as $g(a) = \sum_{i=0}^{d} \alpha_i a^i \in \mathcal{A}$, using the convention $a^0 = 1_{\mathcal{A}}$.) It is known if f is the minimal polynomial of a then the set $\{g \in K[X] | g(a) = 0\}$ is the principal ideal (f) of $K[X]$ generated by f.

If L is an arbitrary extension field of K then the L-space $\mathcal{A}_L = L \otimes_K \mathcal{A}$ can be considered as an L-algebra in a natural way. Multiplication is the K-bilinear extension of

$$\alpha \otimes x \cdot \beta \otimes y = \alpha\beta \otimes xy.$$

$\mathcal{A}$ can be identified with the K-subalgebra $1 \otimes \mathcal{A}$ of $\mathcal{A}_L$. Note that if $a_1, \ldots, a_n$ is a K-basis of $\mathcal{A}$ then $a_1, \ldots, a_n$ is an L-basis of $\mathcal{A}_L$.

$\mathcal{A}$ is called *separable* over K if $\mathcal{A}_L$ is semisimple over any field extension L of K. It turns out that a finite dimensional K-algebra is separable over K if and only if $\mathcal{A}$ is semisimple and the simple components of $\mathrm{C}(\mathcal{A})$ are separable extension fields of K. In particular, every central simple algebra is separable as well as every semisimple algebra over a perfect field.

Ground fields

We are primarily interested in symbolic computations over *global fields.* Global fields are algebraic number fields (finite extensions of the field $\mathbb{Q}$ of the rational numbers) and *global function fields*, that are finitely generated extensions of transcendence degree one over finite fields. Some of our methods have natural extensions to transcendence degree more than one as well. Therefore sometimes we work in the more general setting of *algebraic function fields* over finite fields, i.e., finite extensions of the field $\mathbb{F}_q(X_1, \ldots, X_m)$ of rational functions in m variables over the finite field $\mathbb{F}_q$ consisting of q elements. However, the methods of Chapter 6 rely on arithmetic properties specific to global fields. Global function fields are subject of a beautiful branch of algebraic number theory, called *global class field theory.*

Orders

Let R be a Noetherian integrally closed domain, K be the field of quotients of R and let $\mathcal{A}$ be a finite dimensional algebra over K. An *R-order* in $\mathcal{A}$ is a subring Λ of $\mathcal{A}$ satisfying the following properties:

– Λ is a finitely generated module over R, i.e., there exists a finite set $\{a_1, \ldots, a_N\} \subseteq \mathcal{A}$ such that every element of Λ can be written as a sum $\sum_{i=1}^{N} \alpha_i a_i$ with coefficients $\alpha_i \in R$;

– Λ contains the identity element $1_{\mathcal{A}}$ of $\mathcal{A}$;

– Λ generates $\mathcal{A}$ as a linear space over K.

An important special example of an R-order is the case of *integral structure constants.* Assume that $a_1, \ldots, a_n$ is a basis of $\mathcal{A}$ such that every product $a_i a_j$ written as a linear combination $\sum_{l=1}^{n} \gamma_{ij}^l a_l$ of the basis elements has coefficients $\gamma_{ij}^l \in R$. Then the free R-submodule Λ generated by the basis $a_1, \ldots, a_n$ is a subring of $\mathcal{A}$ and if we in addition assume that $1_{\mathcal{A}}$ has integral coefficients as well (e.g., $1_{\mathcal{A}} = a_1$), then Λ is an R-order. In fact, if R is a principal ideal domain then every R-order is of this form.

An R-order Λ in $\mathcal{A}$ is a *maximal R-order* if it is not a proper subring of any other R-order of $\mathcal{A}$. It is known that in separable K-algebras there exists a maximal order, however, in general it is not unique. For example, for every matrix $a \in \mathrm{GL}_n(\mathbb{Q})$, the ring $a^{-1}\mathrm{M}_n(\mathbb{Z})a$ is a maximal $\mathbb{Z}$-order in the central simple $\mathbb{Q}$-algebra $\mathrm{M}_n(\mathbb{Q})$. (Actually, every maximal $\mathbb{Z}$-order in $\mathrm{M}_n(\mathbb{Q})$ is of this form, however, this fact does not generalize to the case where the ground ring R is not a principal ideal domain). On the other hand, if R is a *Dedekind domain*, i.e. Noetherian integrally closed domain such that every prime ideal of R is maximal and $\mathcal{A}$ is a finite separable extension field of K, then the integral closure Λ of R in $\mathcal{A}$ defined by

$$\Lambda = \{x \in \mathcal{A} | \text{there exists a monic polynomial } f(X) \in R[X] \text{ s. t. } f(x) = 0\}$$

is the a unique maximal R-order in $\mathcal{A}$.

Orders are often used for reducing computation in $\mathcal{A}$ "modulo" certain ideals I of R (computing in the ring $\Lambda/I\Lambda$). In particular, if P is a maximal ideal in the Dedekind domain R and Λ is a maximal R-order in the central simple algebra $\mathcal{A}$, then, the structural invariants of the R/P-algebra $\Lambda/P\Lambda$ do not depend on the choice of Λ. These invariants are called *local invariants* of $\mathcal{A}$ at P. If K is a number field and R is the ring of algebraic integers in K, then the local invariants at the prime ideals of R, together with other invariants corresponding to embeddings of K into $\mathbb{C}$, determine the structure of $\mathcal{A}$ up to isomorphism. Analogous statement holds for the case of global function fields. This fairly nontrivial fact has a beautiful unified formulation in terms *valuations* and *completions*. Phenomena of this flavour, i.e., the possibility to ascertain a "global" property from "local" ones are often referred as *Hasse's principle* for the particular property.

Lie algebras and Cartan subalgebras

We restrict ourselves to finite dimensional Lie algebras. Through the theory of Lie groups, Lie algebras play an important role in the study of certain matrix groups.

A Lie algebra $\mathcal{L}$ over the field K is *nilpotent* if there exists an integer $N > 1$ such that $[\ldots[[x_1, x_2], x_3], \ldots, x_N] = 0$ for arbitrary elements $x_1, \ldots, x_N \in \mathcal{L}$. In every Lie algebra $\mathcal{L}$ there exists a largest nilpotent ideal, called the *nilradical* of $\mathcal{L}$. A Lie algebra $\mathcal{L}$ is *semisimple* if it admits no nontrivial nilpotent ideal. Unlike associative algebras, it is possible that the factoralgebra by the nilradical is not semisimple. However, there exists a smallest ideal, called the *radical* of $\mathcal{L}$, such that the factoralgebra is semisimple. The radical is in fact the largest *solvable* ideal, where solvability is a similar, but weaker property than nilpotence. Like the associative analogue, semisimple Lie algebras are characterized as direct sums of simple Lie algebras. There is a characterization of simple Lie algebras over algebraically closed fields of characteristic zero. The classification of simple Lie algebras over fields of positive characteristic is still a living area of research.

The *normalizer* $\mathrm{N}_{\mathcal{L}}(\mathcal{H})$ of a subalgebra $\mathcal{H}$ of $\mathcal{L}$ is defined as

$$\mathrm{N}_{\mathcal{L}}(\mathcal{H}) = \{x \in \mathcal{L} | [x, y] \in \mathcal{H} \text{ for every } y \in \mathcal{H}\}.$$

$\mathrm{N}_{\mathcal{L}}(\mathcal{H})$ is the largest subalgebra of $\mathcal{L}$ containing $\mathcal{H}$ as an ideal. A subalgebra $\mathcal{H}$ of $\mathcal{L}$ is a *Cartan subalgebra* of $\mathcal{L}$ if $\mathcal{H}$ is nilpotent and $\mathrm{N}_{\mathcal{L}}(\mathcal{H}) = \mathcal{H}$. It is known that if K contains sufficiently many elements (compared to the dimension of $\mathcal{L}$) then $\mathcal{L}$ contains a Cartan subalgebra.

Cartan subalgebras play extremely important role in the theory of Lie algebras, in particular, in classification of simple Lie algebras over fields of characteristic zero.

An interesting example of Lie algebras is $\mathcal{A}_{Lie}$, the Lie algebra of an associative algebra $\mathcal{A}$. This is the same vector space as $\mathcal{A}$ and Lie multiplication is defined by $[x, y] = xy - yx$. The Lie algebras $\mathrm{M}_n(K)_{Lie}$, denoted by $\mathrm{gl}_n(K)$ deserve special interest. Subalgebras of $\mathrm{gl}_n(K)$ are called *matrix Lie algebras.* A *matrix representation* of a Lie algebra $\mathcal{L}$ is a (Lie algebra)-homomorphism $\mathcal{L} \to \mathrm{gl}_n(K)$. The analogue of the regular representation of associative algebras is the *adjoint representation* $\mathrm{ad}_{\mathcal{L}}$ defined as follows. For every $x \in \mathcal{L}$, $\mathrm{ad}_{\mathcal{L}}(x)$ is the K-linear transformation on the vector space $\mathcal{L}$ defined by $\mathrm{ad}_{\mathcal{L}}(x)y = [x, y]$. The kernel of the adjoint representation is the *center* of $\mathcal{L}$, consisting of the elements $x \in \mathcal{L}$ such that $[x, y] = 0$ for every $y \in \mathcal{L}$. Obviously, the center is a nilpotent ideal, therefore the adjoint representation of a semisimple Lie algebra is in fact an embbeding of $\mathcal{L}$ into $\mathrm{gl}_n(K)$. The (associative) subalgebra of $\mathrm{End}_K(\mathcal{L})$ generated by the transformations $\mathrm{ad}_{\mathcal{L}}(x) \quad (x \in \mathcal{L})$ is called the enveloping algebra of $\mathcal{L}$. The simple components of the enveloping algebra of a semisimple Lie algebra $\mathcal{L}$ correspond to the simple components of $\mathcal{L}$, whence the Wedderburn decomposition is a relevant tool to decompose semisimple Lie algebras as well.

1.2 The computational model

Representation of data

We are interested in exact (symbolic) computations over finite fields, algebraic number fields and algebraic function fields (finitely generated transcendental extensions) over finite fields.

To obtain sufficiently general results, we consider nonassociative algebras to be given by a collection of *structure constants.* If $\mathcal{A}$ is a nonassociative algebra over the field K and $a_1, a_2, \ldots, a_n$ is a linear basis of $\mathcal{A}$ over K then multiplication can be described by representing the products $a_i a_j$ as linear combinations of the basis elements a_i

$$a_i a_j = c_{ij1} a_1 + \cdots + c_{ijn} a_n.$$

The coefficients $c_{ijk} \in K$ are called structure constants. We consider algebras to be given as an array of structure constants. Since identities like associativity, (anti-)commutativity, and the Jacobi identity are homogeneous and multilinear, it is sufficient to test the corresponding identities on the basis elements to decide whether $\mathcal{A}$ is an associative, a commutative, or a Lie algebra. An element of $\mathcal{A}$ is represented as the array of its coordinates w.r.t. the basis $a_1, \ldots, a_n$. Substructures (such as subalgebras, ideals, subrings, subspaces) are represented by bases whose elements are given as linear combinations of basis elements of a larger structure.

We use the *dense representation* for elements of K, i.e, every element of K is represented (and inputted) as the array of its coordinates with respect to a basis of K over an appropriate subfield K_0. Note that this is the same as considering K as an algebra over K_0. If K is a finite field $\mathbb{F}_q$ consisting of q elements, then we take $K_0 = \mathbb{F}_p$, the prime field of $\mathbb{F}_q$. If K is an algebraic number field, $K_0 = \mathbb{Q}$ (again the prime field). Algebraic function fields of transcendence degree m over the finite field $\mathbb{F}_q$ are assumed to be given as algebras over $\mathbb{F}_q(X_1, \ldots, X_m)$. Let $d = [K : K_0]$. K can be inputted with structure constants from K_0. Note however that in many cases K is a simple extension of K_0 specified by giving the (monic) minimal polynomial f of a single generating element (primitive element) α over the prime field K_0. This representation can be considered as a special case of the representation with structure constants. The structure constants with respect to the basis $1, \alpha, \alpha^2, \ldots, \alpha^{d-1}$ of K are either zeros and ones, or certain coefficients of f. Even if K is given in this way, we consider f to be given in the dense representation, i.e., an array of d elements from K_0. We can, and often shall, consider n-dimensional K-algebras as algebras of dimension $n \times d$ over K_0. Since the structure constants are assumed to be inputted as arrays of d elements from K_0, polynomial time algorithms for K_0-algebras result in polynomial time algorithms for K-algebras.

A rational number is represented by a not necessarily reduced fraction of two integers. The size of an integer is the number of its binary digits. The size of a rational number r is, however, $\mathrm{size}(p) + \mathrm{size}(q)$, where p/q is the reduced form of r. Modulo p residue classes have size $\lceil \log_2(p+1) \rceil$.

The *height* of a polynomial $0 \neq f \in \mathbb{F}_q[X_1, \ldots, X_m]$ is the maximum of the degrees of f in the variables $X_1, \ldots, X_m$. We use the height as a tool to measure size of objects over the ring $\mathbb{F}_q[X_1, \ldots, X_m]$. Polynomials are considered in the *dense representation*, i.e., if f is of height d then f is viewed as a vector of d^m elements of the ground field, corresponding to the coefficients of the monomials of height at most d. A polynomial $f \in \mathbb{F}_q[X_1, \ldots, X_m]$ of height d has size $\Theta(d^m \log q)$.

A rational function $f \in \mathbb{F}_q(X_1, \ldots, X_m)$ is represented as a quotient of two (not necessarily relatively prime) polynomials. The height of f is, however, the maximum height of the numerator and denominator of its reduced form. The size of a rational function of height d is $\Theta(d^m \log q)$.

The size of compound objects (polynomials, vectors, matrices, etc.) is the sum of the sizes of their components. The height of a compound object over $\mathbb{F}_q(X_1, \ldots, X_m)$ is the maximum height of the components.

Another important way to represent an associative algebra is in the form of a matrix algebra. In this case it suffices to specify a set of matrices which generates the algebra. However, from this representation one can efficiently find a basis of the algebra and structure constants with respect to this basis. In the opposite direction, the regular representation gives an efficient method to obtain a matrix representation from structure constants. Similarly, representation of matrix Lie algebras (subspaces of $\mathrm{M}_n(K)$ closed under the operation $[x, y] = xy - yx$) by structure constants can be efficiently computed. On the other hand, no efficient method is known to find a faithful matrix representation of a Lie algebra. (The celebrated Ado–Iwasawa theorem asserts the existence of such representations. No subexponential bound is known on the degree of the smallest faithful representation in general.) Note however, that in many important cases (such as computing Cartan subalgebras) taking the *adjoint representation* (the analogue of the regular representation of associative algebras) is sufficient for our purposes.

Since integrality generally simplifies our computations, we often compute integral bases of our substructures. If $K = \mathbb{Q}$ or $K = \mathbb{F}_q(X_1, \ldots, X_m)$, K is the field of quotients of a nice factorial domain R, namely that of $R = \mathbb{Z}$ or $R = \mathbb{F}_q[X_1, \ldots, X_m]$, respectively. With some sloppyness we shall call the elements of R *integral* elements. (The terminology is justified by the fact that R is integrally closed in K, i.e., R is the set elements of K integral *over* R.) In these cases by a standard trick we may achieve the situation where the structure constants are integral. If $\delta \in R$ is a common multiple of the denominators of the structure constants of the algebra $\mathcal{A}$ with respect to the basis $a_1, \ldots, a_n$, then the structure constants w.r.t. the basis $\delta a_1, \ldots, \ldots \delta a_n$ are from R.

We shall also work with *R-lattices*, i.e., finitely generated R-submodules of linear K-spaces. Typical examples are free R-lattices given by bases. A set of vectors is a basis of a free lattice if and only if they are linearly independent over K. Note that if $R = \mathbb{Z}$ or $R = \mathbb{F}_q[X]$ then R is a Euclidean domain and every R-lattice is free. Furthermore, from a set of generators of a lattice a basis can be computed in polynomial time by the method of [Fr]. If W is a K-subspace of a K-linear space V given by a basis, then it is very easy to compute an *integral basis* of W, i.e., a basis from the lattice generated by the basis consisting of vectors that have integral coordinates w.r.t. the basis of V.

Basic computations over number fields and finite fields

There are deterministic polynomial time algorithms for the arithmetical operations in K (as well as for polynomial arithmetic over K) if K is a finite field or an algebraic number field. The reader is referred to [Kn] for more details. The basic algorithmic tasks of linear algebra (such as computing ranks, determinants, and solving systems of linear equations) can also be accomplished in deterministic polynomial time. The standard textbook methods (such as Gaussian elimination) use polynomially many arithmetical operations over K. If K is finite, the size of intermediate data cannot explode therefore these methods are directly applicable. In the number field case it will be sufficient to solve linear algebra problems over $\mathbb{Q}$. Polynomial time methods are available to solve systems of linear equations over $\mathbb{Q}$ and over $\mathbb{Z}$ (cf. [Bar], [Ed], [Fr], [KB]).

Basic computations over $\mathbb{F}_q(X_1, \ldots, X_m)$

We summarize here some basic methods for computing over the field $\mathbb{F}_q(X_1, \ldots, X_m)$. We will need bounds on heights in Chapters 3 and 4.

Operations

The arithmetical operations in $\mathbb{F}_q(X_1, \ldots, X_m)$ can be carried out using $(d^m \log q)^{O(1)}$ bit operations, where d is a bound on the height of the operands. Computing a linear combination of l vectors of dimension n has complexity $(d^m ln \log q)^{O(1)}$, where d is a bound on the height of the operands. The complexity of multiplication in an algebra over $\mathbb{F}_q(X_1, \ldots, X_m)$ is $((d\Delta)^m n \log q)^{O(1)}$, where in addition to the preceding notation, Δ is a bound on the height of the structure constants. As we have the bound $(n\Delta)^{O(1)}$ on the height of the output object, we infer that the bit size of the output object is $((n\Delta)^m \log q)^{O(1)}$.

The product and sum of r elements of $\mathbb{F}_q(X_1, \ldots, X_m)$ of heights $d_1, \ldots, d_r$ has height at most $d_1 + \ldots + d_r$, and a similar bound can easily be obtained for linear combinations of vectors. An important case is the addition of *integral* operands, i.e., if the operands are all in $\mathbb{F}_q[X_1, \ldots, X_m]$ or they are vectors with all coordinates in $\mathbb{F}_q[X_1, \ldots, X_m]$. In this case, the height of a sum is bounded by the largest of the heights of the operands.

Height of factors of polynomials

The height of a polynomial over $\mathbb{F}_q(X_1, \ldots, X_m)$ is the maximum height of its coefficients. Let $f, g, h \in \mathbb{F}_q[X_1, \ldots, X_m][X]$ be polynomials such that $f = gh$. Let r resp. s be the smallest indices such that the coefficients of X^r resp. X^s have the largest height among the coefficients of g, h, resp. Then all other summands in the coefficient of X^{r+s} have height less than the height of the product of the coefficients of X^r and X^s. We infer that in the ring of the polynomials with integral coefficients the height of a factor of a polynomial f is not greater than the height of f.

Specializations

We often need "sufficiently many" relatively prime maximal ideals of $\mathbb{F}_q[X_1, \ldots, X_m]$ such that the residue class fields are small finite fields. We have q^m *specializations* over the ground field: maximal ideals of type $((X_1 - c_1), \ldots, (X_m - c_m))$, $c_i \in \mathbb{F}_q$. In this case the residue class field is $\mathbb{F}_q$. In some cases (when q is small) we shall work with a suitably chosen finite extension of $\mathbb{F}_q$.

Linear algebra problems

Determinant Let the height of a matrix M from $\mathrm{M}_n(\mathbb{F}_q(X_1, \ldots, X_m))$ be Δ. If M is integral (i.e., $M \in \mathrm{M}_n(\mathbb{F}_q[X_1, \ldots, X_m])$) then the height of its determinant is bounded by $n\Delta$. In the general case we have the bound $n^3\Delta$. (The numerators become of height at most $n^2\Delta$ when we clear the denominators.) We can compute the determinant of an integral matrix via Chinese Remaindering: we specialize the matrix at $(n\Delta + 1)^m$ places from I^m, where I is a subset of cardinality $n\Delta + 1$ of the ground field $\mathbb{F}_q$ (or of an extension of degree $\lceil log_q(n\Delta + 1) \rceil$ if q is small), compute the determinant in the residue class fields and then using Lagrange interpolation compute in each step $(n\Delta + 1)^{m-i}$ interpolating polynomials of maximum degree at most $n\Delta$ in $\mathbb{F}_q[X_1, \ldots, X_i]$.

Matrix inversion, nonsingular systems of linear equations With the aid of determinants we can readily compute the inverse of a matrix from $\mathrm{M}_n(\mathbb{F}_q[X_1, \ldots, X_m])$ of height Δ. In this (not necessarily reduced) representation, the elements of the inverse matrix have numerators of height at most $(n - 1)\Delta$ and a common denominator of height at most $n\Delta$. Thus, as a simple observation, we have that the height of the inverse of an integral matrix will be at most $n\Delta$; for an arbitrary element of $\mathrm{M}_n(\mathbb{F}_q(X_1, \ldots, X_m))$ we have a factor n^3 instead of n. Similarly, determinants allow us to use Cramer's rule for solving nonsingular systems of linear equations. For the height of the solution, we have bounds similar to the bounds for inverse matrices.

Linear independence of vectors, rank of matrices. If we have l integral vectors (i.e., the components are in $\mathbb{F}_q[X_1, \ldots, X_m]$) of length n and height at most Δ, then the rank of the matrix consisting of these vectors is at most $\min(n, l)$ and the height of the subdeterminants is at most $\min(n, l)\Delta$, hence we can test their independence via

$(1 + \min(n, l)\Delta)^m$ specializations. The case of matrices whose components are general rational functions can be reduced to the case of integral matrices of the same rank and of height $\min(n, l)\Delta$ by clearing denominators.

Homogeneous systems of linear equations Given a homogeneous system of l linear equations for n variables, the solution is the kernel V of the $n \times l$ matrix of the system and our objective is to compute a basis of V. We can assume that the coefficients are integral (we can readily obtain an equivalent system with coefficients from $\mathbb{F}_q[X_1, \ldots, X_m]$ of height at most nl times larger than the height of the original coefficients) and compute an integral basis of the solution. Let us denote the rank of the matrix by r. We can obtain a basis of the solution space in the standard way. This involves solving $n - r$ nonsingular systems of linear equations in r variables each. Multiplication with the determinant provides integral solutions. In this way we obtain an integral basis of the solution space of height at most $r\Delta$, where Δ is a bound on the height of the coefficients.

Alternatively, if we already know a bound Γ on the height of an integral basis of the solution space ($\min(n, l)\Delta$ in general), then we can solve a homogeneous system of linear equations over the ground field $\mathbb{F}_q$ consisting of at most $(\Gamma + \Delta)^m$ equations for Γ^m variables (i.e., the system describing that all the coefficients vanish), and from a solution choose a maximal independent set over $\mathbb{F}_q[X_1, \ldots, X_m]$.

Elementary computations in algebras

The identity element of $\mathcal{A}$, if exists, can obviously be obtained as a solution of a system of linear equations. The same holds for certain substructures such as the center and centralizer. Bases of subalgebras, left ideals, right ideals, two-sided ideals generated by a finite subset $X \subset \mathcal{A}$, as well as structure constants for subalgebras and factoralgebras can also be computed in straightforward ways. The minimal polynomial of an element $x \in \mathcal{A}$ can be computed using the standard method. There are obvious polynomial bounds on the sizes of these objects in algebras over number fields as well as on their heights in the function field case.

f-algorithms and ff-algorithms

It will be handy to use some conventions introduced by Rónyai ([Ró3, Ró4]). Some of our methods rely on solutions of subproblems not known to have deterministic polynomial time algorithms. These subproblems are finding the prime factorization of an integer and factoring polynomials over a finite prime field. An *ff-algorithm* is a deterministic method if it is allowed to call oracles for these two problems. Similarly an *f-algorithm* is a deterministic method which is allowed to call an oracle for factoring polynomials over finite fields. In both cases the cost of a call is the size of the input of the call.

The use of f-algorithms is convenient because of the fact that the monic polynomials $g \in K[X]$ dividing a polynomial $f \in K[X]$ are in one-to-one correspondence with the ideals in the factoralgebra $K[X]/(f)$, whence factoring polynomials over K is in fact a subcase of, say, finding the maximal ideals of commutative K-algebras.

The first deterministic polynomial time algorithm for factoring polynomials over $\mathbb{Q}$ was proposed in the seminal paper [LLL]. This result was later extended to arbitrary number fields in [Ch], [Gri], [La], and [Len]. For factoring polynomials over finite fields a deterministic method was given by Berlekamp in [Ber1, Ber2]. The method is based on a deterministic polynomial time reduction to factor polynomials that split into linear factors over the prime field, and a brute force search method for solving the latter special case. The time complexity of this algorithm is polynomial in the parameters $p, \log_p q$ and $\deg(f)$, where $f \in \mathbb{F}_q[x]$ is the polynomial to be factored and the characteristic of $\mathbb{F}_q$ is the prime p. Note that the input size is in fact $\Theta(\log q \deg f)$, therefore the running time of the method is not polynomial in the input size. In [Ber3], Berlekamp proposed a randomized (*Las Vegas*) factoring algorithm that runs in time polynomial in the input size. (In contrast to *Monte Carlo* methods, Las Vegas methods never give incorrect answer.) It follows that a polynomial time f-algorithm can be replaced with a polynomial Las Vegas method.

For factoring integers no polynomial time methods are known, neither deterministic nor randomized. This problem is widely believed to be difficult. We will use ff-algorithms for some problems related to simple algebras over number fields.

1.3 Previous results

In this section we give a short summary of the most important results in the area of computing the structure of algebras. For more bibliographic details the reader is referred to the survey paper [Ró4]. We continue to use the term algebra for a finite dimensional associative algebra over the field K (given by structure constants). The ground field is always either a finite field or an algebraic number field.

The radical

The first method for computing the radical of algebras over fields of zero characteristic is due to Dickson [Di]. The method is based on a characterization via a system of linear equations.

In [Ró2], L. Rónyai proposed an analogous, although much more sophisticated characterization of the radical of algebras over finite prime fields. This characterization results in a deterministic polynomial time algorithm for computing the radical of algebras over finite fields. Note that Eberly extended the characterization to algebras over arbitrary finite fields.

Wedderburn decomposition

The first efficient algorithm for computing the minimal ideals of semisimple algebras over finite fields and algebraic number fields was given by K. Friedl (cf. [FR]). The method is an iteration based on factoring minimal polynomials of a basis of the center over certain *extensions* of the ground field obtained from earlier steps of iteration. In the number field case, the method is a deterministic polynomial one, while in the finite case the method is a polynomial time f-algorithm.

Eberly in [Eb2] presented a polynomial time Las Vegas algorithm which avoids iteration. The key idea is that (under the reasonable assumption that the ground field has sufficiently many elements) a random element of a commutative semisimple algebra is in fact a generating element. (We note that this method can be derandomized using the techniques of Chapter 5.) In the same paper, a deterministic (and parallelizable) reduction to factoring minimal polynomials of the basis elements of the center is also presented.

In [GMT], Gianni, Miller and Trager outline a method, based on lifting primitive idempotents modulo an appropriate small prime, for computing the Wedderburn decomposition of a commutative semisimple algebra over $\mathbb{Q}$. The running time is exponential in the dimension. The authors claim that a combination with lattice basis reduction techniques of [LLL] leads to a polynomial time method. In fact, our method of Chapter 4 could be applied to algebras over $\mathbb{Q}$.

Decomposition of simple algebras

In this subsection we denote by $\mathcal{A}$ a central simple algebra of dimension n^2 over the field K.

In the case when K is finite, by a theorem of Wedderburn, $\mathcal{A}$ is isomorphic to $\mathrm{M}_n(K)$. In [Ró2], L. Rónyai proposed a polynomial time f-algorithm for computing such an isomorphism.

The problem of decomposition of simple algebras over number fields appears to be much more difficult. In [Ró1], Rónyai gives a Las Vegas polynomial time reduction from the *quadratic residuosity* problem to computing the index of central simple algebras of dimension 4 (so-called *quaternion algebras*) over $\mathbb{Q}$. Note that the result is conditional on the Generalized Riemann Hypothesis (GRH for short). (For generalizations of the Riemann Hypothesis and their significance in computational number theory the reader is referred to Bach [Bach].) The quadratic residuosity problem, formulated by Goldwasser and Micali in [GM], is to decide whether a number is quadratic residue modulo a squarefree number, and is believed to be difficult. It is also shown in [Ró1] that finding a zero divisor in a quaternion algebra over $\mathbb{Q}$ is (again under GRH) at least as hard as finding solutions of quadratic congruences $x^2 \equiv a \pmod{n}$ (taking a square root of a if exists) modulo a squarefree number n, which is, up to a Las Vegas polynomial time reduction (see [Ra] or [GM]) is as hard as factoring n. This fact justifies the use of ff-algorithms to solve related problems.

On the other hand, Rónyai proved [Ró3] that the decision problem related to computing the index of the central simple algebra $\mathcal{A}$ over the number field K is in $NP \cap coNP$. In fact, the existence of a maximal order with short description and verification is proved and the result is combined with a technique, based on Hasse's principle to compute the index from a maximal order. For testing maximality of an order a polynomial time ff-algorithm is used.

An easier task is to compute an isomorphism $\mathcal{A}_L \cong \mathrm{M}_n(L)$ for an appropriate extension L of K. Using again the technique of random elements, Eberly in [Eb1] and [Eb3] presents a Las Vegas polynomial time method to construct such an extension L together with an isomorphism $\mathcal{A}_L \cong \mathrm{M}_n(L)$. He applies this to compute isomorphism $\mathcal{A}_{\mathbb{R}} \cong \mathrm{M}_n(\mathbb{R})$ or $\mathcal{A}_{\mathbb{R}} \cong \mathrm{M}_{\frac{n}{2}}(\mathbb{H})$ for embeddings $K \to \mathbb{R}$. Here, $\mathbb{H}$ stands for the skewfield of the Hamiltionian quaternions. Note that this method can be derandomized using the results in [Ró5] or techniques of Chapter 5.

Chapter 2

Testing membership in abelian matrix groups over number fields

In this brief chapter we present an application of algebra decompositions over number fields to a basic problem related to matrix groups. This material has been published as a part of the paper [BBCIL], joint work with László Babai, Robert Beals, Jin-yi Cai, and Eugene M. Luks. Let K be a number field, n and r be positive integers, and $h, g_1, \ldots, g_r \in \mathrm{GL}_n(K)$ be invertible n by n matrices over K. We assume the dense representation, i.e., matrices are inputted as arrays of n^2 elements from K, where elements of K are represented as arrays of $d = [K : \mathbb{Q}]$ rational numbers and K is given by structure constants (or by the minimal polynomial of a primitive element) over $\mathbb{Q}$. The *membership problem* is the problem of deciding whether h is in the subgroup G of $\mathrm{GL}_n(K)$ generated by $g_1, \ldots, g_r$. The *constructive membership problem* is, in addition to testing membership, to express h in terms of the generators in the case when h is in the group G. Note that the membership problem is in general undecidable for $n \geq 4$ (see [Mi]). We restrict ourselves to the abelian case, i.e., we assume that the matrices $h, g_1, \ldots, g_r \in \mathrm{GL}_n(K)$ are pairwise commuting. (This condition can be efficiently tested in the straightforward way.)

The *constructive membership problem for abelian matrix groups* is to test whether the equation

$$(*) \qquad g_1^{x_1} \cdots g_r^{x_r} = h$$

admits an integer solution $(x_1, \ldots, x_r) \in \mathbb{Z}^r$, and if it does, find such a solution.

We present a deterministic polynomial time algorithm for this problem. Our method is based on a reduction to the case $n = 1$, which was recently solved by G. Ge in [Ge1, Ge2].

Ge's theorem. *Given an algebraic number field K and nonzero elements $\alpha_1, \ldots, \alpha_r \in K$, one can in polynomial time compute a basis of the lattice consisting of the solutions*

$(x_1, \ldots, x_r) \in \mathbb{Z}^r$ *to the equation*

$$1 = \alpha_1^{x_1} \cdots \alpha_r^{x_r}.$$

Note that the analogous problem for commutative semigroups (where the matrices $a_1, \ldots, a_r$ are not necessarily regular but the exponents $x_1, \ldots, x_r$ are required to be non-negative) was solved for the special case $r = 2$ in [CLZ]. For generalizations of the problem the reader is referred to the paper [BBCIL]. More recent developments for the membership problem in matrix groups can be found in [Bea].

2.1 The algorithm

First we observe that it is sufficient to find bases of lattices given as solutions to equations of the form

$$(**) \qquad g_1^{x_1} \cdots g_r^{x_r} = \mathrm{Id}_n,$$

a special case of $(*)$ taking $h = \mathrm{Id}_n$. Indeed, we take $g_0 = h^{-1}$, introduce a new variable x_0, and find a basis of the lattice L of the solutions to the equation

$$g_0^{x_0} g_1^{x_1} \cdots g_r^{x_r} = \mathrm{Id}_n$$

in $r + 1$ variables. An element of L with first coordinate $x_0 = 1$ can be found by solving a linear equation over $\mathbb{Z}$.

Let $\mathcal{A} \leq \mathrm{M}_n(K)$ be the subalgebra of $\mathrm{M}_n(K)$ generated by the matrices $g_1, \ldots, g_r$. Obviously, $\mathcal{A}$ is commutative and, since g_1 is invertible, contains the identity matrix Id_n. We can compute a basis of $\mathcal{A}$ and the corresponding structure constants in polynomial time. We use Dickson's method [Di] to compute the radical $\mathrm{Rad}(\mathcal{A})$, and then the method [FR] to compute the simple components $\overline{\mathcal{A}_1}, \ldots, \overline{\mathcal{A}_s}$ of the factoralgebra $\overline{\mathcal{A}} = \mathcal{A}/\mathrm{Rad}(\mathcal{A})$. For every $j \in \{1, \ldots, s\}$, we compute the maximal ideal

$$\sum_{l \neq j} \overline{\mathcal{A}_l}$$

of $\overline{\mathcal{A}}$ complementary to $\overline{\mathcal{A}_j}$ and hence the natural homomorphism $\phi_j : \mathcal{A} \to \overline{\mathcal{A}_j}$. Since $\overline{\mathcal{A}_j}$ is simple, we can apply Ge's method to compute bases of the lattices L_j given by

$$L_j = \{(x_1, \ldots, x_r) \in \mathbb{Z}^r | \prod_{i=1}^{r} \phi_j(g_i)^{x_i} = 1_{\overline{\mathcal{A}_j}}\}.$$

A basis

$$\underline{b}_1 = (b_{11}, b_{12}, \ldots, b_{1r}), \ldots, \underline{b}_t = (b_{t1}, b_{t2} \ldots, b_{tr})$$

of the intersection

$$L = \bigcap_{j=1}^{s} L_j$$

can then be found in polynomial time via solving a system of linear equations over $\mathbb{Z}$. Since for the natural homomorphism $\phi : \mathcal{A} \to \overline{\mathcal{A}}$ we have

$$\phi = \bigoplus_{j=1}^{s} \phi_j.$$

L is in fact the lattice of the solutions to the equation

$$\prod_{i=1}^{s} \phi(g_i)^{x_i} = 1_{\overline{\mathcal{A}}}.$$

Obviously, L contains the solutions of $(**)$. Therefore it is sufficient to look for the solutions of $(**)$ in terms of the vectors $\underline{b}_1, \ldots, \underline{b}_t$. In other words, our task is to construct a basis of solutions $(y_1, \ldots, y_t) \in \mathbb{Z}^t$ to

$$\prod_{i=1}^{r} g_i^{\sum_{j=1}^{t} b_{ji} y_j} = \mathrm{Id}_n,$$

which can also be written as

$$(***) \qquad \prod_{j=1}^{t} {g'_j}^{y_j} = \mathrm{Id}_n,$$

where for every $j \in \{1, \ldots, t\}$, the matrix g'_j is defined as

$$g'_j = \prod_{i=1}^{r} g_i^{b_{ji}}.$$

Since $\underline{b}_1, \ldots, \underline{b}_r$ are in L, we have $\phi(g'_j) = 1_{\overline{\mathcal{A}}}$, whence $g'_j - \mathrm{Id}_n \in \mathrm{Rad}(\mathcal{A})$ for every $j = 1, \ldots, t$.

In an algebra $\mathcal{A}$ with identity the elements $u = 1 + v$ where $v \in \mathrm{Rad}(\mathcal{A})$ form a subgroup $\mathrm{U}(\mathcal{A})$ (called the *unipotent radical*) in the multiplicative group $\mathcal{A}^*$ of units (invertible elements). For an element $u = 1 + v \in 1 + \mathrm{Rad}(\mathcal{A})$, we define the logarithm $\log u \in \mathrm{Rad}(\mathcal{A})$ to be the sum

$$\log u = \sum_{i=1}^{\infty} \frac{(-1)^{i-1}}{i} v^i.$$

(Note that this sum has only at most $\dim_K \mathrm{Rad}(\mathcal{A})$ nonzero terms.) The logarithm-map $\log :\ \mathrm{U}(\mathcal{A}) \to \mathrm{Rad}(\mathcal{A})$ is invertible, the inverse is

$$\exp(v) = \sum_{i=0}^{\infty} \frac{1}{i!} v^i$$

for $v \in \mathrm{Rad}(\mathcal{A})$. In addition, if $\mathcal{A}$ is commutative (as in our case) then these maps are group isomorphisms $(\mathrm{U}(\mathcal{A}), \cdot) \cong (\mathrm{Rad}(\mathcal{A}), +)$. It follows that equation $(***)$ is equivalent to

$$\sum_{j=1}^{t} y_j \log g'_j = 0.$$

Expanding this equation w.r.t. matrix entries, we obtain a system of n^2 homogeneous linear equations with coefficients from K, which is, after further expansion, equivalent to a system of $n^2 \times [K : \mathbb{Q}]$ homogeneous linear equations with coefficients from $\mathbb{Q}$. After clearing denominators, we obtain a system with coefficients from $\mathbb{Z}$, which can be solved in polynomial time.

We have proved the following

Theorem 2.1.1 *The constructive membership problem for commutative matrix groups over number fields can be solved in deterministic polynomial time.* $\square$

Chapter 3

Computing the radical

The material presented in this chapter is a combination obtained from parts of the papers [IRSz] (joint work with Lajos Rónyai and Ágnes Szántó) and [CIW] (joint work with Arjeh M. Cohen and David B. Wales). In [Di], Dickson gave a nice characterization of the radical of a matrix algebra $\mathcal{A} \leq \mathrm{M}_n(K)$, where $\mathrm{char}K = 0$. Namely, $\mathrm{Rad}(\mathcal{A})$ is the largest left (right, or two-sided) ideal L of $\mathcal{A}$ such that the trace of every element of L is zero. This characterization leads to an efficient computation of $\mathrm{Rad}(\mathcal{A})$. It can be obtained as the solution space of a system of homogeneous linear equations.

If K is of positive characteristic p then the trace of a matrix algebra $\mathcal{A} \leq \mathrm{M}_n(K)$ can vanish even if $\mathcal{A}$ is semisimple. For the case $K = \mathbb{F}_p$, Rónyai introduced in [Ró2] a new linear function $\mathrm{Tr}' : \mathcal{A} \to K$ when the ordinary trace is identically zero on $\mathcal{A}$. This new function can still vanish on $\mathcal{A}$, but then a further linear function Tr'' can be introduced, and so on. However, if $\mathcal{A}$ is not nilpotent then this procedure terminates in at most $\lceil \log_p n \rceil$ rounds. This leads to a method analogous to Dickson's algorithm for computing the radical. A decreasing sequence of ideals of $\mathcal{A}$ can be computed using solutions of systems of linear equations. The sequence collapses to the radical in at most $\lceil \log_p n \rceil$ steps. The construction of the functions is based on integral lifts of matrices.

Eberly extended Rónyai's results to matrix algebras over arbitrary finite fields [Eb1]. The construction of Eberly's functions is still based on lifting matrices to characteristic zero. The new functions are semilinear rather than linear if the ground field is greater than the prime field.

Here we present a construction based on the paper [CIW] that works in matrix algebras over an arbitrary field of positive characteristic. The functions are defined as certain coefficients of the characteristic polynomial.

Section 3.1 is devoted to the definitions and basic properties of the generalized trace functions and a characterization of the radical that extend the above mentioned results of

Rónyai and Eberly. In Section 3.2, we relate our functions to those used in [Ró2], [Eb1] and [IRSz]. In the zero characteristic case, the values of the trace function on a basis of $\mathcal{A}$ are known to determine the composition factors of the underlying $\mathcal{A}$-module. In Section 3.3, we give a generalization of this fact. In Section 3.4, based on work [IRSz], we present an algorithm to compute the radical of algebras over finitely generated pure transcendental extensions of finite fields.

Throughout this chapter, $\mathcal{A}$ denotes a finite dimensional associative algebra over the field K. By U, V, W, etc. we denote finite dimensional $\mathcal{A}$-modules. For standard facts and definitions related to modules and representations the reader is referred to textbooks, e.g., [Pie]. To avoid confusion, we fix here some minor details of the terminology. The $\mathcal{A}$-module $\{0\}$ is called the *trivial* $\mathcal{A}$-module. An $\mathcal{A}$-module Z is a *zero module* if $az = 0$ for every $a \in \mathcal{A}$ and $z \in Z$. An $\mathcal{A}$-module V is called *simple* or *irreducible* if V is not a zero module and V admits exactly two submodules: $\{0\}$ and V. The composition factors of a nontrivial module V are either simple modules or one-dimensional zero modules. We shall refer to the composition factors of V that are simple modules as the *nonzero composition factors* of V.

3.1 Trace functions and the radical

Let V be a finite dimensional $\mathcal{A}$-module. For $a \in \mathcal{A}$ we denote the action of a on V by a_V. This means, that $a_V \in \mathrm{End}_K V$ is the linear transformation $v \mapsto av$. The characteristic polynomial $\chi_{V,a}(X)$ of the action of a on V is simply the characteristic polynomial $\chi_{a_V}(X) = \det(a_V - X \cdot \mathrm{Id}_V)$ of the linear transformation a_V. For our purposes it appears to be more convenient to use the variant

$$\widetilde{\chi}_{V,a}(X) = \det(X \cdot a_V + \mathrm{Id}_V) = X^{\dim_K V} \chi_{V,a}(-1/X)$$

of the characteristic polynomial. For an integer $s > 0$ we define the s'th trace $\mathrm{Tr}_V(s,a)$ of the action a on V as the s'th coefficient of the polynomial $\widetilde{\chi}_{V,a}(X)$ (considered as a formal power series in X):

$$\widetilde{\chi}_{V,a}(X) = 1 + \sum_{s=1}^{\infty} \mathrm{Tr}_V(s,a) X^s.$$

Obviously, $\mathrm{Tr}_V(s,a) = 0$ for $s > \dim V$, while $\mathrm{Tr}_V(1,a), \mathrm{Tr}_V(2,a), \ldots, \mathrm{Tr}_V(\dim V, a)$ are up to sign the coefficients of the characteristic polynomial:

$$(-1)^{\dim V} \chi_{V,a}(X) = X^{\dim V} + \sum_{s=1}^{\dim V} (-1)^s \mathrm{Tr}_V(s,a) X^{\dim V - s}.$$

$\mathrm{Tr}_V(1,\cdot)$ coincides with the ordinary trace function $\mathrm{Tr}_V(\cdot)$, therefore it is linear on $\mathcal{A}$. Note that, for general s, $\mathrm{Tr}_V(s,a)$ is the trace of the (diagonal) action of a on the s'th exterior power of V.

If U is a submodule of V, by choosing a basis appropriately, the corresponding matrix representation has a block-upper triangular form. It is obvious that for every $a \in \mathcal{A}$,

$$\widetilde{\chi}_{V,a}(X) = \widetilde{\chi}_{U,a}(X)\widetilde{\chi}_{V/U,a}(X).$$

It follows that if two $\mathcal{A}$-modules V and U have the same composition factors (counted with multiplicities), then $\mathrm{Tr}_U(s,a) = \mathrm{Tr}_V(s,a)$ for every positive integer s and $a \in \mathcal{A}$. Note that if Z is a zero module, then $\widetilde{\chi}_{Z,a}(X)$ is identically 1 for every $a \in \mathcal{A}$. Therefore if the nonzero composition factors of W and V coincide then $\mathrm{Tr}_W(s,a) = \mathrm{Tr}_V(s,a)$.

If $\mathcal{A}$ is semisimple then Wedderburn's theorems say that $\mathcal{A}$ is a direct sum of full matrix rings over division algebras. The irreducible nonzero modules are the natural modules on exactly one of these matrix rings. If $\mathcal{A}$ is not semisimple, then $\mathrm{Rad}(\mathcal{A})$ is the intersection of the annihilators of all the irreducible modules. As a consequence, the trace functions are in fact defined on $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$, i.e., $\mathrm{Tr}_V(s,a+b) = \mathrm{Tr}_V(s,a)$ for every $\mathcal{A}$-module V, integer $s > 0$, $a \in \mathcal{A}$, and $b \in \mathrm{Rad}(\mathcal{A})$. In particular, the fuctions $\mathrm{Tr}_V(s,\cdot)$ vanish on $\mathrm{Rad}(\mathcal{A})$. If V is a faithful $\mathcal{A}$-module and the functions $\mathrm{Tr}_V(s,\cdot)$ are identically zero on $\mathcal{A}$ for every integer $s > 0$, then every $a \in \mathcal{A}$ is nilpotent, whence $\mathcal{A}$ is nilpotent. It follows that if V is faithful, and L is a left (or right) ideal in $\mathcal{A}$ such that the functions $\mathrm{Tr}_V(s,\cdot)$ are identically zero on L for every integer $s > 0$, then L is a nilpotent one sided ideal in $\mathcal{A}$, whence $L \subseteq \mathrm{Rad}(\mathcal{A})$. As a consequence, (assuming again that V is faithful), we have

$$\mathrm{Rad}(\mathcal{A}) = \{a \in \mathcal{A} | \mathrm{Tr}_V(s,a) = \mathrm{Tr}_V(s,ba) = 0 \text{ for every } s > 0 \text{ and } b \in \mathcal{A}\}.$$

If K is of characteristic zero then Dickson's classical result [Di] is that

$$\mathrm{Rad}(\mathcal{A}) = \{a \in \mathcal{A} | \mathrm{Tr}_V(1,a) = \mathrm{Tr}_V(1,ba) = 0 \text{ for every } b \in \mathcal{A}\},$$

As $\mathrm{Tr}_V(1,\cdot)$ is linear on $\mathcal{A}$, this characterization leads to an efficient algorithm for $\mathrm{Rad}(\mathcal{A})$: it can be obtained by solving of a system of homogeneous linear equations.

Our aim is to obtain a similar result in positive characteristic. From now on we assume that K is of positive characteristic p. First we observe that there is an obvious sufficient condition for that trace functions of low index vanish on $\mathcal{A}$. In that case, some higher trace function will be semilinear.

Proposition 3.1.1 *Assume that K is of positive characteristic p and the multiplicities of the nonzero composition factors of V are all divisible by p^j for an integer $j > 0$. Then*

$\mathrm{Tr}_V(s,a) = 0$ *for every* $a \in \mathcal{A}$ *and positive integer* s *such that* $p^j \not| s$. *Furthermore, the function* $\mathrm{Tr}_V(p^j,\cdot)$ *is* p^j*-semilinear on* $\mathcal{A}$*, i.e., for every* $\alpha,\beta \in K$ *and* $a,b \in \mathcal{A}$*, we have*

$$\mathrm{Tr}_V(p^j, \alpha a + \beta b) = \alpha^{p^j}\mathrm{Tr}_V(p^j,a) + \beta^{p^j}\mathrm{Tr}_V(p^j,b).$$

Proof Let $V_1,\ldots,V_h$ be representatives of the isomorphism classes of the nonzero irreducible $\mathcal{A}$-modules occurring in a composition series of V with multiplicities $m_1p^j,\ldots,m_hp^j$. Let $W = V_1^{m_1} \oplus \ldots \oplus V_h^{m_h}$, the formal direct sum of m_1 copies of V_1, m_2 copies of $V_2,\ldots,$ and m_h copies of V_h. Obviously, for every $a \in \mathcal{A}$ we have $\widetilde{\chi}_{V,a}(X) = \widetilde{\chi}_{W,a}(X)^{p^j}$, whence $\mathrm{Tr}_V(s,a)$ is zero for every positive integer s not divisible by p^j, $\mathrm{Tr}_V(p^j,a) = \mathrm{Tr}_W(1,a)^{p^j}$, and semilinearity of $\mathrm{Tr}_V(p^j,\cdot)$ follows from the linearity of the ordinary trace function $\mathrm{Tr}_W(1,\cdot) = \mathrm{Tr}_W(\cdot)$. $\square$

If K is algebraically closed, then the converse also holds.

Proposition 3.1.2 *Assume that* K *is an algebraically closed field of positive characteristic* p*. Let* j *be a positive integer and assume that* $\mathrm{Tr}_V(p^i,a) = 0$ *for every* $a \in \mathcal{A}$ *and* $0 \le i < j$*. Then the multiplicities of the nonzero composition factors of* V *are all divisible by* p^j*.*

Proof Since the values $\mathrm{Tr}_V(p^i,a)$ depend only on $a+\mathrm{Rad}(\mathcal{A})$ and the nonzero composition factors of V, we may without loss of generality assume that both $\mathcal{A}$ and V are semisimple. Assume that V contains an irreducible constituent W with multiplicity mp^t, where $t < j$ and p does not divide m. Since K is algebraically closed, the action of $\mathcal{A}$ on W is isomorphic to the full matrix ring $\mathrm{M}_d(K)$ (where $d = \dim_K W$) and there is a primitive idempotent e in $\mathcal{A}$ such that e has matrix

$$\begin{pmatrix} 1 & 0 & \cdot & \cdot & 0 \\ 0 & 0 & \cdot & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & \cdot & 0 \end{pmatrix}$$

in its action on W with respect to some basis. Since e is a primitive idempotent, it acts as zero on the irreducible $\mathcal{A}$-modules not isomorphic to W. It follows that

$$\widetilde{\chi}_{V,e}(X) = \widetilde{\chi}_{W,e}(X)^{mp^t} = (1+X)^{mp^t} = (1+mX+\ldots)^{p^t} = 1 + mX^{p^t} + \ldots,$$

in other words, $\mathrm{Tr}(p^t,e) = m\cdot 1_K \neq 0$. $\square$

We define a sequence of subsets $\mathcal{A} = \mathcal{A}_{V,0} \supseteq \mathcal{A}_{V,1} \supseteq \mathcal{A}_{V,2}\ldots$ inductively by

$$\mathcal{A}_{V,j+1} = \{a \in \mathcal{A}_{V,j} | \mathrm{Tr}_V(p^j,a) = \mathrm{Tr}_V(p^j,ba) = 0 \text{ for every } b \in \mathcal{A}_{V,j}\}.$$

Obviously, for every j, $\mathcal{A}_{V,j} \supseteq \mathrm{Rad}(\mathcal{A})$. If the ground field is algebraically closed, propositions 3.1.1 and 3.1.2 give a characterization of the subsets $\mathcal{A}_{V,j}$. It generalizes the interpretation of the radical $\mathrm{Rad}(\mathcal{A})$ of $\mathcal{A}$ as the annihilator of all the simple $\mathcal{A}$-modules:

$$\mathrm{Rad}(\mathcal{A}) = \{a \in \mathcal{A} \mid aL = 0 \text{ for every simple } \mathcal{A}\text{-module } L\}.$$

Proposition 3.1.3 *Assume that K is an algebraically closed field of characteristic $p > 0$. Let the composition factors of V be $V_1, \ldots, V_r$ with multiplicities $m_1, \ldots, m_r$. Then for every integer $j \geq 0$, $\mathcal{A}_{V,j}$ equals the annihilator*

$$I_{V,j} = \{a \in \mathcal{A} \mid aV_s = 0 \text{ for every } s \text{ such that } p^j \not| m_s\}.$$

As a consequence, for every j, $\mathcal{A}_{V,j}$ is an ideal in $\mathcal{A}$.

Proof Again, w.l.o.g., we may assume that $\mathcal{A}$ is semisimple. For every $j \geq 0$, the annihilator $I_{V,j}$ is the intersection of the kernels of the actions of $\mathcal{A}$ on the simple modules V_s such that $p^j \not| m_s$, therefore $I_{V,j}$ is an ideal. Proposition 3.1.1 implies that for every j, $I_{V,j} \subseteq \mathcal{A}_{V,j}$. We prove that $\mathcal{A}_{V,j} = I_{V,j}$ by induction on j. The assertion is trivial for $j = 0$. Assume that for some $j \geq 0$ we have $\mathcal{A}_{V,j} = I_{V,j}$. By Proposition 3.1.1, $\mathrm{Tr}(p^j, \cdot)$ is semilinear on $I_{V,j}$, whence $\mathcal{A}_{V,j+1}$ is a left ideal in $I_{V,j}$, therefore it is a left ideal of $\mathcal{A}$ as well. In particular, $\mathcal{A}_{V,j+1}$ is a subalgebra, and from Proposition 3.1.2 we infer that the multiplicities of the nonzero composition factors of V as an $\mathcal{A}_{V,j+1}$-module are all divisible by p^{j+1}. On the other hand, since $\mathcal{A}_{V,j+1}$ is a left ideal, the nonzero simple $\mathcal{A}_{V,j+1}$-modules are simple $\mathcal{A}$-modules as well, giving that $\mathcal{A}_{V,j+1} \subseteq I_{V,j+1}$. □

Now we are ready to prove the most important properties of the trace functions.

Theorem 3.1.4 *Let K be an arbitrary field of characteristic $p > 0$. For every integer $j \geq 0$, we have*

(i) $\mathcal{A}_{V,j}$ is an ideal in $\mathcal{A}$ containing $\mathrm{Rad}(\mathcal{A})$,

(ii) $\mathrm{Tr}_V(s, a) = 0$ for every $a \in \mathcal{A}_{V,j}$ and positive integer s such that $p^j \not| s$,

(iii) The function $\mathrm{Tr}_V(p^j, \cdot)$ is p^j-semilinear on $\mathcal{A}_{V,j}$ in the sense that for every $a, b \in \mathcal{A}_{V,j}$ and $\alpha, \beta \in K$, we have

$$\mathrm{Tr}_V(p^j, \alpha a + \beta b) = \alpha^{p^j} \mathrm{Tr}_V(p^j, a) + \beta^{p^j} \mathrm{Tr}_V(p^j, b)$$

(iv) If V is a faithful $\mathcal{A}$-module and $p^j > \dim_K V$, then $\mathcal{A}_{V,j} = \mathrm{Rad}(\mathcal{A})$.

Proof Let $\bar{K}$ be an algebraic closure of K. Then $\mathcal{A}$ can be identified with the subalgebra $1 \otimes \mathcal{A}$ of $\overline{\mathcal{A}} = \bar{K} \otimes_K \mathcal{A}$. Also, if V is a K-module then $\overline{V} = \bar{K} \otimes V$ can considered to be an $\overline{\mathcal{A}}$-module in a natural way. (The multiplication is the K-bilinear extension of $(\alpha \otimes a) \cdot (\beta \otimes v) := (\alpha\beta) \otimes (av)$.) If we fix a K-basis $v_1, \ldots, v_m$ of V then $1 \otimes v_1, \ldots, 1 \otimes v_m$ is a $\bar{K}$-basis of $\overline{V}$, and the matrix of the action of $a \in \mathcal{A}$ on V with respect to the basis $v_1, \ldots, v_m$ is also the matrix of the action of $1 \otimes a$ on $\overline{V}$ with respect to the basis $1 \otimes v_1, \ldots, 1 \otimes v_m$. It follows that the characteristic polynomials χ, their variants $\widetilde{\chi}$, and therefore our trace functions remain the same on $\mathcal{A}$. We have also $\mathrm{Rad}(\mathcal{A}) = \mathcal{A} \cap \mathrm{Rad}(\overline{\mathcal{A}})$. (Altough there are examples where $\dim_K \mathrm{Rad}(\mathcal{A}) < \dim_{\bar{K}} \mathrm{Rad}(\overline{\mathcal{A}})$.)

The propositions proved above state that all the assertions hold for $\overline{\mathcal{A}}$ and $\overline{V}$ in place of $\mathcal{A}$ and V, respectively. Obviously, it is enough to prove the equalities

$$\mathcal{A}_{V,j} = \mathcal{A} \cap \overline{\mathcal{A}}_{\overline{V},j}, \quad (j = 0, 1, 2, \ldots).$$

But this fact can be proved by an easy induction on j, using the semilinearity of the function $\mathrm{Tr}_{\overline{V}}(p^j, \cdot)$ on $\overline{\mathcal{A}}_{\overline{V},j}$. $\square$

3.2 Computing trace functions via lifting

In this section we relate our functions to those used in [Ró2], [Eb1], or [IRSz]. We need the following elementary combinatorial fact.

Lemma 3.2.1 *Let m and i be positive integers and consider the permutation group G on $\Omega = \{1, \ldots, m^i\}$ generated by the cyclic permutation $\alpha = (12 \cdots m^i)$. Let $0 < j \leq i$. Then the size of an orbit of G acting on the j-element subsets of Ω is divisible by m^{i-j+1}.*

Proof A subgroup $H \leq G$ of order r has m^i/r orbits of size r in Ω. The subsets of Ω stabilized by H are unions of some of these orbits and hence their sizes are divisible by r. Now if H is the stabilizer of a j-element subset of Ω and $|H| = r$, then $r \mid j$ and the size of the orbit of G (among the j-element subsets) containing this set is m^i/r. Suppose indirectly that m^i/r is not divisible by m^{i-j+1}. Then there exists a prime power factor q of m, such that q^{i-j+1} does not divide m^i/r, whence $q^{j-1} \neq r$ and $q^{j-1} \mid r \mid j$. We have that $2^{j-1} \leq q^{j-1} < r \leq j$, which is impossible if $j \geq 1$. $\square$

Let R be an arbitrary commutative ring with identity and P be an ideal in R. We shall use the customary congruence-notation, i.e, for $x, y \in R$ we write $x \equiv y \pmod{P}$ if $x - y \in P$.

Proposition 3.2.2 *Let n be a positive integer, R be a commutative ring with identity and $a, b \in \mathrm{M}_n(R)$. Let m be a positive rational integer and P be an ideal of R containing $m1_R$. Suppose that $a \equiv b \pmod{P\mathrm{M}_n(R)}$. Then*

$$\mathrm{Tr}(a^{m^j}) \equiv \mathrm{Tr}(b^{m^j}) \pmod{P^{j+1}}$$

for all nonnegative integers j.

Proof Let $d = b - a$, then $d \in P\mathrm{M}_n(R)$.

$$b^{m^j} - a^{m^j} = (a+d)^{m^j} - a^{m^j} = \sum_{\emptyset \neq \sigma \subseteq \{1,\ldots,m^j\}} u_1^\sigma u_2^\sigma \cdots u_{m^j}^\sigma,$$

where $u_r^\sigma = d$, if $r \in \sigma$; $u_r^\sigma = a$ otherwise. Since $\mathrm{Tr}(vw) = \mathrm{Tr}(wv)$, the traces of the terms in one orbit of the group G generated by $\pi = (12\cdots m^i)$ are all the same. The trace of a term containing i d's is in P^i. If $i > j$, then we obtain immediately that the contribution of the orbit is in P^{j+1}. If $0 < i \leq j$, then sum of the traces of the terms in a G-orbit is in $m^{j-i+1}P^i \subseteq P^{j+1}$ by the preceding lemma. □

We use this in the situation where $m = p$, a rational prime, P is a maximal ideal in R above p, and $K = R/P$. We have that, for a matrix a over R, the residue class of $\mathrm{Tr}(a^{p^j})$ modulo P^{j+1} depends only on the residue classes of the entries of the matrix a modulo P. In other words the function $a \mapsto \mathrm{Tr}(a^{p^j}) + P^{j+1}$ is in fact a function $\mathrm{M}_n(K) \to R/P^{j+1}$. We use a notation analogous to that we introduced in the preceding section. For $a \in \mathrm{M}_n(R)$, $\mathrm{Tr}(s, a)$ stands for the sth coefficient of the polynomial $\widetilde{\chi}_a(X) = \det(\mathrm{Id}_n + Xa) \in R[X]$. (Since we work with a fixed module $V = R^n$, we omit the subscript V.)

Proposition 3.2.3 *Let n be a positive integer, R be a commutative ring with identity element and p be a rational prime. Let j be a positive integer and P an ideal of R containing $p1_R$. Assume that $a \in \mathrm{M}_n(R)$ is an $n \times n$ matrix with entries from R such that*

$$\mathrm{Tr}(p^i, a^h) \equiv 0 \pmod{P}$$

for every positive integer h and nonnegative integer $i < j$. Then

$$\mathrm{Tr}(a^{p^j}) \equiv (-1)^{p^j-1} p^j \mathrm{Tr}(p^j, a) \pmod{P^{j+1}}.$$

Proof The congruence holds trivially if $j = 0$. Assume that $j > 0$. Let $\bar{a} \in \mathrm{M}_n(R/P)$ be the residue of the matrix a modulo P. By Theorem 3.1.4, applied to the matrix algebra generated by $\bar{a}$, the assumption on a is equivalent to that the coefficients $\mathrm{Tr}(s, a)$ of the characteristic polynomial $\chi_a(X)$ are in P whenever $p^j \not| s$.

Recall that the companion matrix C_f of a polynomial $f(X) \in R[X]$ of the form $X^n - \alpha_1 X^{n-1} - \ldots - \alpha_n$ is

$$\begin{pmatrix} 0 & 0 & 0 & \cdot & \cdot & \cdot & \alpha_n \\ 1 & 0 & 0 & \cdot & \cdot & \cdot & \alpha_{n-1} \\ 0 & 1 & 0 & \cdot & \cdot & \cdot & \alpha_{n-1} \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & \cdot & \cdot & \cdot & 0 & 1 & \alpha_1 \end{pmatrix}.$$

Moreover the (usual) characteristic polynomial of C_f is exactly f. Newton's identities

$$\mathrm{Tr}(a^s) + (-1)^s s\mathrm{Tr}(s,a) + \sum_{r=1}^{s-1}(-1)^r \mathrm{Tr}(a^{s-r})\mathrm{Tr}(r,a) = 0 \quad (s = 1,\ldots,n),$$

$$\mathrm{Tr}(a^s) + \sum_{r=1}^{n}(-1)^r \mathrm{Tr}(a^{s-r})\mathrm{Tr}(r,a) = 0 \quad (s > n)$$

imply that the traces $\mathrm{Tr}(a^s)$ are uniquely determined by the coefficients $\mathrm{Tr}(s,a)$ of the characteristic polynomial of a. We infer that both sides of the congruence to be proved remain the same if we replace a with the companion matrix of $\chi_a(X)$. Therefore it is enough to prove the congruence for matrices of the form above such that $\alpha_s \in P$ for every s such that $p^j \not| s$. By Proposition 3.2.2, we may further assume that $0 = \alpha_s = \pm\mathrm{Tr}(s,a)$ for every s such that $p^j \not| s$. In the case $p^j > n$ both sides are zero. If $p^j \leq n$, the p^jth Newton identity immediately gives $\mathrm{Tr}(a^{p^j}) + (-1)^{p^j} p^j \mathrm{Tr}(p^j,a) = 0$. $\square$

Assume that R is an integral domain of characteristic zero such that $P = pR$ is a maximal ideal in R and $K \cong R/pR$. Obviously, if $K = \mathbb{F}_p$, we can take $R = \mathbb{Z}$. A construction for $q = p^h$, $K = \mathbb{F}_q[X_1,\ldots,X_m]$ (possibly with $m = 0$, when $K = \mathbb{F}_q$), can be the local ring $R = \mathbb{Z}[X][X_1,\ldots,X_m]_{(p)}/g(X)$, where $g(X) \in \mathbb{Z}[X]$ is a monic lift of a monic irreducible polynomial in $\mathbb{F}_p[X]$. Assume that $\mathcal{A}$ is a subalgebra of $\mathrm{M}_n(K)$. As in the preceding paragraph, we can define the ideals $\mathcal{A} = \mathcal{A}_0 \supseteq \mathcal{A}_1 \supseteq \mathcal{A}_2 \supseteq \ldots$ of $\mathcal{A}$ by

$$\mathcal{A}_{j+1} = \{a \in \mathcal{A}_j | \mathrm{Tr}(p^j,a) \equiv \mathrm{Tr}(p^j,ba) \equiv 0 \pmod{P} \text{ for every } b \in \mathcal{A}_j\}.$$

In [Ró2], [Eb1], and [IRSz], the fuctions $T^{(j)} : \mathcal{A}_j \to K$ defined by the congruence

$$p^j T^{(j)}(a) \equiv \mathrm{Tr}(a^{p^j}) \pmod{P^{j+1}}$$

are used, where R and $P \lhd R$ are as described above. The last proposition asserts that these functions up to sign coincide with our trace functions.

3.3 Trace functions and composition factors

We know that if the nonzero composition factors of two $\mathcal{A}$-modules W and V coincide, meaning that they are the same as multisets, then $\mathrm{Tr}_W(s,a) = \mathrm{Tr}_V(s,a)$ for every $a \in \mathcal{A}$. Also, if K is of characteristic zero then the converse holds (see [Di]) in the following strong sense: If B is a basis of $\mathcal{A}$ and $\mathrm{Tr}_W(1,a) = \mathrm{Tr}_V(1,a)$ for every $a \in B$ then the nonzero composition factors of W and V coincide. We generalize this to positive characteristic. We will need the following additivity property.

Lemma 3.3.1 *Assume that K is of characteristic $p > 0$. Let U be a submodule of V. Suppose that $\mathrm{Tr}_U(p^i,a) = 0$ for every $1 \leq i < j$ and $a \in \mathcal{A}$ or $\mathrm{Tr}_{V/U}(p^i,a) = 0$ for every $1 \leq i < j$ and $a \in \mathcal{A}$. Then*

$$\mathrm{Tr}_V(p^j,a) = \mathrm{Tr}_U(p^j,a) + \mathrm{Tr}_{V/U}(p^j,a)$$

for every $a \in \mathcal{A}$.

Proof Obviously

$$\widetilde{\chi}_{V,a}(X) = \widetilde{\chi}_{U,a}(X)\widetilde{\chi}_{V/U,a}(X),$$

whence

$$\mathrm{Tr}_V(p^j,a) = \mathrm{Tr}_U(p^j,a) + \mathrm{Tr}_{V/U}(p^j,a) + \sum_{r,s\geq 1, r+s=p^j} \mathrm{Tr}_U(r,a)\mathrm{Tr}_{V/U}(s,a).$$

By Theorem 3.1.4 applied to U or V/U, we have $\mathrm{Tr}_U(s,a) = 0$ for $1 \leq s < p^j$ or $\mathrm{Tr}_{V/U}(s,a) = 0$ for $1 \leq s < p^j$, giving the statement. $\square$

Theorem 3.3.2 *Let K be of characteristic $p > 0$. Suppose $B \subseteq \mathcal{A}$ is a subset of $\mathcal{A}$ such that $B \cup \mathrm{Rad}(\mathcal{A})$ is a K-linear generating set of $\mathcal{A}$ (e.g., B induces a basis of $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$). For the $\mathcal{A}$-modules U and V the following assertions are equivalent.*

(i) U and V have isomorphic nonzero composition factors (counted with multiplicities);

(ii) $\mathrm{Tr}_U(s,a) = \mathrm{Tr}_V(s,a)$ for every $a \in \mathcal{A}$ and positive integer s;

(iii) $\mathrm{Tr}_U(p^j,a) = \mathrm{Tr}_V(p^j,a)$ for every $a \in \mathcal{A}$ and nonnegative integer j;

(iv) $\mathrm{Tr}_U(p^j,a) = \mathrm{Tr}_V(p^j,a)$ for every $a \in B$ and $j \leq \max\{\log_p \dim_K U, \log_p \dim_K V\}$.

Proof Assertion (ii) can be reformulated as follows.

(ii′) $\widetilde{\chi}_{U,a}(X) = \widetilde{\chi}_{V,a}(X)$ for every $a \in \mathcal{A}$

As the terms composition factors and trace functions are defined modulo $\mathrm{Rad}(\mathcal{A})$, we may assume that $\mathcal{A}$ is semisimple and the modules U and V are unital. In that case, condition (i) is equivalent to

(i′) U and V are isomorphic $\mathcal{A}$-modules.

Implications (i) $\Rightarrow$ (ii′) $\Rightarrow$ (iii) $\Rightarrow$ (iv) are obvious. First we prove the implication (iv) $\Rightarrow$ (ii′). Assume that (iv) holds. Let

$$N = \max\{\lceil \log_p \dim_K U\rceil, \lceil \log_p \dim_K V\rceil\}$$

and $W = U^{p^N-1} \oplus V$ (direct sum of $p^N - 1$ copies of U and one of V). We work in $U \oplus W$. We have

$$\widetilde{\chi}_{U\oplus W,a}(X) = (\widetilde{\chi}_{U,a}(X))^{p^N} \widetilde{\chi}_{V,a}(X) \equiv \widetilde{\chi}_{V,a}(X) \pmod{X^{p^N}}.$$

In other words, $\mathrm{Tr}_{U\oplus W}(s, a) = \mathrm{Tr}_V(s, a)$ for every $a \in \mathcal{A}$ and integer $0 < s < p^N$.

We prove by induction on j that $\mathrm{Tr}_W(p^j, \cdot)$ is identically zero on $\mathcal{A}$ if $j < N$. For $j = 0$, it follows from the linearity of the ordinary trace. Assume that $0 < j < N$ and the functions $\mathrm{Tr}_W(p^i, \cdot)$ are identically zero on $\mathcal{A}$ for $i < j$. Lemma 3.3.1 implies that

$$\mathrm{Tr}_W(p^j, a) = \mathrm{Tr}_{U\oplus W}(p^j, a) - \mathrm{Tr}_U(p^j, a).$$

We have seen that the expression on the right hand side is $\mathrm{Tr}_V(p^j, a) - \mathrm{Tr}_U(p^j, a)$, which is zero for every $a \in B$ by condition (iv). By Theorem 3.1.4, the function $\mathrm{Tr}_W(p^j, \cdot)$ must be identically zero on $\mathcal{A}$.

We have proved that the functions $\mathrm{Tr}_W(p^j, \cdot)$ $0 \le j < N$ vanish on $\mathcal{A}$. Theorem 3.1.4 implies that the same holds for the functions $\mathrm{Tr}_W(s, \cdot)$ for every $1 \le s < p^N$. In other words, for every $a \in \mathcal{A}$,

$$\widetilde{\chi}_{W,a}(X) \equiv 1 \pmod{X^{p^N}},$$

whence

$$\widetilde{\chi}_{U\oplus W,a}(X) \equiv \widetilde{\chi}_{U,a}(X) \pmod{X^{p^N}}.$$

Combining with the congruence

$$\widetilde{\chi}_{U\oplus W,a}(X) \equiv \widetilde{\chi}_{V,a}(X) \pmod{X^{p^N}}$$

and looking at the degrees we infer $\widetilde{\chi}_{V,a}(X) = \widetilde{\chi}_{U,a}(X)$.

We prove implication (ii′) $\Rightarrow$ (i′) by induction on $\dim_K U$. Assume that (ii′) holds. Note that if U is not the trivial module, then there is an s such that $\mathrm{Tr}_U(s, \cdot)$ is not identically zero on $\mathcal{A}$, as witnessed by the identity of $\mathcal{A}$. We immediately see that U is the trivial

$\mathcal{A}$-module iff V is the trivial $\mathcal{A}$-module as all traces are 0. Let s be the smallest positive integer such that $\mathrm{Tr}_U(s,\cdot)$ is not identically zero on $\mathcal{A}$. By Theorem 3.1.4 we have $s = p^j$ for some nonnegative integer j and $\mathrm{Tr}_U(s,\cdot) = \mathrm{Tr}_V(s,\cdot)$ is semilinear on $\mathcal{A}$. Since $\mathcal{A}$ is the sum of its minimal left ideals, there exists a minimal left ideal L such that $\mathrm{Tr}_U(s,\cdot) = \mathrm{Tr}_V(s,\cdot)$ is not identically zero on L. It is known that L is zero on simple $\mathcal{A}$-modules not isomorphic to L (as a left $\mathcal{A}$-module). Therefore there exist submodules $U_L \leq U$ and $V_L \leq V$ isomorphic to L. Since for every $a \in \mathcal{A}$ we have

$$\widetilde{\chi}_{U/U_L,a}(X) = \widetilde{\chi}_{U,a}(X)/\widetilde{\chi}_{L,a}(X) \quad \text{and} \quad \widetilde{\chi}_{V/V_L,a}(X) = \widetilde{\chi}_{V,a}(X)/\widetilde{\chi}_{L,a}(X),$$

assumption (ii$'$) is inherited by U/U_L and V/V_L, and we can use induction in the case when $U > U_L$. □

Combining with the characteristic 0 case we have the following result.

Corollary 3.3.3 *Let K be an arbitrary field and $B \subseteq \mathcal{A}$ such that $B \cup \mathrm{Rad}(A)$ is a K-linear generating set of $\mathcal{A}$ (e.g., B induces a K-basis of $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$). For the $\mathcal{A}$-modules U and V the following assertions are equivalent.*

(i) U and V have isomorphic composition factors (counted with multiplicities);

(ii) $\chi_{U,a}(X) = \chi_{V,a}(X)$ for every $a \in \mathcal{A}$;

(iii) $\chi_{U,a}(X) = \chi_{V,a}(X)$ for every $a \in B$.

□

3.4 Algorithms

In this section K is a field of characteristic $p > 0$. The results of Section 3.1 suggest a method for computing $\mathrm{Rad}(\mathcal{A})$ based on solving p^j-semilinear equations. The following simple example demonstrates that the problem of computing the radical of finite dimensional algebras over K involves taking p'th roots of elements of K.

Example Let K be a field of characteristic p, $\alpha \in K$, and $\mathcal{A} = K[X]/(X^p - \alpha)$. It is known that the polynomial $X^p - \alpha$ is either irreducible in $K[X]$ or has a unique root $\alpha^{\frac{1}{p}}$ in K. In the former case $\mathrm{Rad}(\mathcal{A})$ is zero, while in the latter the inverse image of $\mathrm{Rad}(\mathcal{A})$ is the ideal I in $\mathcal{A}[X]$ generated by $X - \alpha^{\frac{1}{p}}$. Using standard polynomial arithmetic over K, from $\mathrm{Rad}(\mathcal{A})$ we can compute the unique monic generator $X - \alpha^{\frac{1}{p}}$ of I. □

The following construction, due to Fröhlich and Shepardson [FSh], demonstrates that it is impossible to take pth roots (therefore computing the radicals of finite dimensional algebras) by algorithms using merely the field operations.

Example Let L be an enumerable but not recursive set and f be an enumeration of L without repetition, i.e., a recursive function $\mathbb{N} \to \mathbb{N}$, such that f is a bijection $\mathbb{N} \to L$. Standard examples exist, such as an appropriate encoding of the halting problem of Turing machines, or that of Hilbert's tenth problem. We set

$$\mathbb{F}_0 = \mathbb{F}_q(Y_i \mid i \in \mathbb{N}), \quad R = \mathbb{F}_0[X_i \mid i \in \mathbb{N}],$$

and $K = R/I$, where I is the ideal of R generated by all of the polynomials $X_i^p - Y_{f(i)}$ $(i \in \mathbb{N})$. It is straightforward to see that K is the extension field of F_0 obtained by adjoining p'th roots of all $Y_{f(i)}$ $(i \in \mathbb{N})$.

An element of K is represented by a fraction of polynomials in (a finite number of) the variables X_i, Y_j. The field operations as well as equality tests can be carried out using straightforward calculation in subfields generated by the symbols involved in the operands. Obviously, the a variable Y_i has a p'th root in K if and only if $i \in L$, therefore membership test in L can be reduced to the problem of deciding whether an element of K has a pth root in K, whence the latter problem is also undecidable. $\Box$

To get around this, we assume that K is a finite extension of the subfield

$$K^p = \{x^p | x \in K\}$$

and we can effectively compute a K^p-basis of this extension. Note that since $\alpha \mapsto \alpha^p$ is an isomorphism $K \cong K^p$, this implies that for every positive integer j, K is finite over K^{p^j}, too. More precisely, we have $[K : K^{p^j}] = [K : K^p]^j$. Also, we can effectively compute a K^{p^j}-basis of K.

In the next proposition we summarize the main computational tasks one encounters when trying to apply the characterization given in Theorem 3.1.4 to compute the radical. As we work in a quite general setting, the issue of size (of field elements obtained during the computation) is not addressed here.

Proposition 3.4.1 *Let n be a positive integer and suppose that the field K of positive characteristic p is a finite extension of $K^{p^{\lfloor \log_p(n+1) \rfloor}}$. Let $\mathcal{A}$ be an n-dimensional associative algebra over K given by structure constants. Then the problem of finding a basis of the radical of $\mathcal{A}$ can be reduced with $(n[K : K^{p^{\lfloor \log_p(n+1) \rfloor}}])^{O(1)}$ arithmetical operations in K to the following computational subtasks.*
For each j, $0 \le j \le \lfloor \log_p(n+1) \rfloor$,

(i) computing a basis of K over the subfield K^{p^j};

(ii) computing the value of $\mathrm{Tr}(p^j, x)$ of at most $n(n+1)$ at most $(n+1)$-dimensional matrices over K;

(iii) computing the coordinates of at most $n(n+1)$ elements of K in the basis computed in (i) over the subfield K^{p^j}.

(iv) solving a system of linear equations over the field K^{p^j}. A system has at most $(n+1)[K : K^{p^j}]$ equations and at most n variables;

(v) computing the p^j'th root $x^{\frac{1}{p^j}}$ of at most n^2 elements $x \in K^{p^j}$.

Proof We work with the regular representation of the algebra $\mathcal{A}$ or the regular representation of the Dorroh extension of $\mathcal{A}$, if $\mathcal{A}$ has no identity element. In either case, we obtain a faithful representation of $\mathcal{A}$ with $\dim_K \mathcal{A} = \nu$. ($\nu = n$ or $\nu = n+1$.) We identify $\mathcal{A}$ with its image at this representation. We use the notation of Section 3.1 omitting the subscript V. Let $\{a_1, \ldots, a_n\}$ be a basis of $\mathcal{A}$ over K. Let $a_0 := \mathrm{Id}_\nu$. Note that we need this element only if $\mathcal{A}$ itself has no identity element.

To prove the proposition, we outline a computation for producing a basis over K of $\mathcal{A}_j$ from a basis of $\mathcal{A}_{j-1}$. Theorem 3.1.4 shows that $\mathrm{Rad}(\mathcal{A}) = \mathcal{A}_{\lceil \log_p \nu \rceil}$. For an integer $j \geq 0$, suppose that we are already given a basis $\{b_1, \ldots, b_h\}$ of $\mathcal{A}_{j-1}$ over K. By the p^j-semilinearity of $\mathrm{Tr}(p^j, \cdot)$, for the coefficients α_l of elements $\sum \alpha_l b_l \in \mathcal{A}_j$ we have the following system of equations:

$$\sum_{l=1}^{h} \alpha_l^{p^j} \mathrm{Tr}(p^j, b_l a_i) = 0, \;\; i = 0, \ldots, n.$$

We compute first the values $\mathrm{Tr}(p^j, b_l a_i)$, and their coordinates t_{jlis} over the subfield K^{p^j}. Next we solve the system of linear equations over K^{p^j}

$$\sum_{l=1}^{h} \beta_l t_{jlis} = 0, \;\; i = 0, \ldots, n, \;\; s = 1, \ldots, [K : K^{p^j}].$$

As a solution, we obtain a K^{p^j}-basis $\{(\beta_{11}, \ldots, \beta_{1h}), \ldots, (\beta_{r1}, \ldots, \beta_{rh})\}$ of the solution space $V \subset (K^{p^j})^h$. The Frobenius map $\Phi_j \colon (K)^h \to (K^{p^j})^h$ is an isomorphism, consequently the inverse images of the elements of this basis will form a basis of $\mathcal{A}_j$ over K. $\square$

We intend to use the algorithm outlined in the proposition for the field $K = \mathbb{F}_q(X_1, \ldots, X_m)$. For the index of the subfield $K^{p^{\lfloor \log_p (n+1) \rfloor}}$ we have

$$[K : K^{p^{\lfloor \log_p (n+1) \rfloor}}] = (p^{\lfloor \log_p (n+1) \rfloor})^m = O(n^m).$$

Also we can efficiently compute the values of the inverse of the Frobenius map. The characteristic polynomial of an $n \times n$ matrix a over K can also be computed in time $(n + \mathrm{height}(a))^{mO(1)}$. The only serious problem could be explosure of data sizes. We have to solve about $\log_p n$ systems of linear equations in an iteration, where the coefficients of the j'th system depend on the solution of the previous system. We need a bound on the height of a basis of the ideals $\mathcal{A}_j$. For a step of the iteration, we have the following.

Lemma 3.4.2 *Let $\mathcal{A}$ be an n-dimensional algebra over the field $K = \mathbb{F}_q(X_1, \ldots, X_m)$. Assume that the structure constants have numerators of height at most Δ and a common denominator in $\mathbb{F}_q[X_1, \ldots, X_m]$. Assume further that for a $0 \leq j \leq \lfloor log_p n \rfloor$, the ideal $\mathcal{A}_{j-1}$ has an integral basis (in the sense that the coordinates of the basis elements of $\mathcal{A}_{j-1}$ with respect to the input basis of $\mathcal{A}$ are from $\mathbb{F}_q[X_1, \ldots, X_m]$) of height at most Γ. Then the ideal $\mathcal{A}_j$ has an integral basis of height at most $n(\Gamma + 2\Delta)$.*
In particular, if the algebra $\mathcal{A}$ is primary, i.e., if $\mathcal{A}/\mathrm{Rad}(\mathcal{A}))$ is simple, then $\mathrm{Rad}(\mathcal{A})$ has an integral basis of height at most $2n\Delta$.

Proof The endomorphism $\Phi_j \colon K \to K^{p^j}$ induces an automorphism of the finite field $\mathbb{F}_q$ and maps X_i to $Y_i = X_i^{p^j}$, therefore

$$K^{p^j} = \mathbb{F}_q(Y_1, \ldots, Y_m) = \mathbb{F}_q(X_1^{p^j}, \ldots, X_m^{p^j})$$

and

$$\{X_1^{i_1} \cdots X_m^{i_m} \mid 0 \leq i_1, \ldots, i_m < p^j\}$$

is a basis of K over K^{p^j}. Let us use the notation of the proof of Proposition 3.4.1. The elements $b_l a_i$ expressed in terms of a_i have a common denominator and numerators of height at most $\Gamma + \Delta$. We infer a similar statement for the entries of the $\nu \times \nu$ matrices $b_l a_i$. The bound on the height of the numerators becomes $\Gamma + 2\Delta$. It follows that the expressions $\mathrm{Tr}(p^j, (b_l a_i))$ have numerators of height at most $p^j(\Gamma + 2\Delta)$ and a common denominator. By clearing the denominator we may assume that $t_{jils} \in \mathbb{F}_q[Y_1, \ldots, Y_m]$ of degree at most $\Gamma + 2\Delta$ in each Y_r. Our system of linear equations has rank at most n. This implies that there exists an integral basis of the solution space of degree at most $n(\Gamma + 2\Delta)$ (in $Y_1, \ldots, Y_m$). By substituting X_i in Y_i we obtain a basis of the ideal $\mathcal{A}_j$. If $\mathcal{A}/\mathrm{Rad}(\mathcal{A})$ is simple, then $\mathcal{A}_j = \mathcal{A}_{j-1}$ except for at most one j, therefore the height can increase in at most one step. $\square$

Unfortunately, from this lemma immediately we could only obtain a quasi-polynomial $n^{c \log n}\Delta$ bound on the radical. The next two statements allow us to obtain manageable bounds on the height of the radical. Also, they turn out to be useful later when we consider the decomposition of semisimple algebras. Our argument is essentially a version of the technique of Eberly [Eb1, Eb2, Eb3] based on *splitting elements.*

Lemma 3.4.3 *Let $\mathcal{A}$ be an n-dimensional algebra over the field $K = \mathbb{F}_q(X_1, \ldots, X_m)$. Assume that the structure constants are from $\mathbb{F}_q[X_1, \ldots, X_m]$ and their heights are bounded by Δ. Then for any ideal I of $\mathcal{A}$ properly containing $\mathrm{Rad}(\mathcal{A})$ there exists an ideal I' having an integral basis of height at most $2(n\lfloor 2log_q n\rfloor + (n+1)\Delta)$, such that I' is contained in I but not contained in $\mathrm{Rad}(\mathcal{A})$.*

Proof First we explain the construction over an arbitrary (but sufficiently large) field K. We consider the semisimple algebra $\bar{\mathcal{A}} = \mathcal{A}/\mathrm{Rad}(\mathcal{A})$ over K. Let $\bar{I}$ be the image of an ideal $I \neq \mathcal{A}$ under the natural map and let $\bar{J}$ be the complement ideal of $\bar{I}$ in $\bar{\mathcal{A}}$. Let us examine the right action of the elements of $\mathcal{A}$ on $\bar{I}$ and $\bar{J}$. The resultant of the characteristic polynomials of the right action on $\bar{I}$ resp. $\bar{J}$ of a generic element $\sum_{i=1}^{n} z_i a_i$ of $\mathcal{A}$ is a not identically zero polynomial in $K[z_1, \ldots, z_n]$, and its degree in each variable is at most $\dim_K \bar{I} \dim_K \bar{J} \leq n^2/4$. The sparse zeros lemma (cf. [Sch] or Section 5.1 in this work) implies that if the cardinality of a subset $L \subseteq K$ is greater than $n^2/4$ then there exists an element of the form $a = \sum_{i=1}^{n} l_i a_i$, all $l_i \in L$, such that the characteristic polynomials f resp. g of the right action of a on $\bar{I}$ resp. $\bar{J}$ are relatively prime. We now substitute a into g. The element $g(a)$ annihilates $\bar{J}$, but does not annihilate $\bar{I}$. We conclude that $g(a) \in I \setminus \mathrm{Rad}(\mathcal{A})$.

Returning to the situation at hand $K = \mathbb{F}_q(X_1, \ldots, X_m)$, we select L from elements of small height:

$$L = \{h \in \mathbb{F}_q[X_1] \mid \deg h \leq \lfloor 2 \log_q n \rfloor\}$$

will suffice. The preceding argument demonstrates the existence of an integral element a of height at most $\lfloor 2log_q n\rfloor$ with the property $g(a) \in I \setminus \mathrm{Rad}(\mathcal{A})$. The characteristic polynomial g of the right action of a on $\bar{J}$ is a monic polynomial dividing the monic characteristic polynomial with integral coefficients of the right action on $\mathcal{A}$ and the heights of the coefficients of the latter polynomial are bounded by $n(\lfloor 2log_q n\rfloor + \Delta)$. We infer that g has also integral coefficients of height at most $n(\lfloor 2log_q n\rfloor + \Delta)$ and we can bound the height of $g(a)$ by $2n(\lfloor 2log_q n\rfloor + \Delta)$. The ideal I' generated by $g(a)$ has an integral basis chosen from the elements $\{a_i g(a) a_j \mid i, j = 1, \ldots, n\}$ of height $2(n\lfloor 2log_q n\rfloor + (n+1)\Delta)$. □

We can now prove a polynomial bound on the heights of the ideals over the radical:

Proposition 3.4.4 *Let $\mathcal{A}$ be an n-dimensional algebra over the field $K = \mathbb{F}_q(X_1, \ldots, X_m)$. Assume that the structure constants are integral (i.e., they are from $\mathbb{F}_q[X_1, \ldots, X_m]$) and their heights are limited by Δ. Then any ideal of $\mathcal{A}$ containing $\mathrm{Rad}(\mathcal{A})$ has an integral basis of height $O(n^3(\lfloor 2log_q n\rfloor + \Delta))$.*

Proof Let $I_1, \ldots, I_r$ be the minimal elements of the set of the ideals of $\mathcal{A}$ containing $\mathrm{Rad}(\mathcal{A})$. Let $I'_1, \ldots, I'_r$ be ideals corresponding to them according to Lemma 3.4.3 with integral bases of height at most

$$\Gamma = 2(n\lfloor 2log_q n\rfloor + (n+1)\Delta).$$

The minimality of the ideals I_j implies that $I'_j + \mathrm{Rad}(\mathcal{A}) = I_j$. Then the structure constants of I'_j have a common denominator and numerators of height at most $\Delta + (n+1)\Gamma$. (This can be seen by writing up systems of linear equations for the structure constants of the ideals I'_j. The coefficients of the variables are of height at most Γ, while for the heights of the inhomogeneous parts – corresponding to products of basis elements of I'_j – the bound is $2\Gamma + \Delta$.) We have $\mathrm{Rad}(I'_j) = I'_j \cap \mathrm{Rad}(\mathcal{A})$ and Noether's isomorphism implies that

$$(*) \qquad I'_j/\mathrm{Rad}(I'_j) \cong (I'_j + \mathrm{Rad}(\mathcal{A}))/\mathrm{Rad}(\mathcal{A}) = I_j/\mathrm{Rad}(\mathcal{A}).$$

We infer that $I'_j/\mathrm{Rad}(I'_j)$ is a simple algebra and we can apply Lemma 3.4.2. For the height of the radical of I'_j we have the bound $2n(\Delta + (n+1)\Gamma)$ with respect to the above basis of I'_j and the bound $\Sigma = 2n(\Delta + (n+1)\Gamma) + \Gamma$ with respect to the original basis of $\mathcal{A}$. Also infer that

$$J = \mathrm{Rad}(I'_1) + \ldots + \mathrm{Rad}(I'_r) \subseteq \mathrm{Rad}(\mathcal{A})$$

is an ideal of $\mathcal{A}$ having an integral basis of height at most Σ. Moreover $J \cap I'_j = \mathrm{Rad}(I'_j)$ holds for $1 \le j \le r$. Now an application of Noether's isomorphism theorem and $(*)$ gives that the algebras $(I'_j + J)/J$ are simple. For the ideal $I = I'_1 + \ldots + I'_r$ this entails that I/J is semisimple as a sum of the simple algebras $(I'_1 + J)/J, \ldots, (I'_r + J)/J$, whence $\mathrm{Rad}(I) = J$. The residue class algebra

$$\mathcal{A}/I = (\mathrm{Rad}(\mathcal{A}) + I)/I \cong \mathrm{Rad}(\mathcal{A})/(I \cap \mathrm{Rad}(\mathcal{A})$$

is a radical algebra, therefore if we assume that $\mathcal{A}$ has an identity element, then $I = \mathcal{A}$ and then the height of $\mathrm{Rad}(\mathcal{A})$ is at most Σ. If $\mathcal{A}$ has no identity element, then we can compute the radical of $\mathcal{A}$ in its Dorroh extension $\tilde{\mathcal{A}}$, since $\mathrm{Rad}(\tilde{\mathcal{A}}) = \mathrm{Rad}(\mathcal{A})$. The latter relation follows from $\tilde{\mathcal{A}}/\mathcal{A} \cong K$. We have obtained that $\mathrm{Rad}(\mathcal{A})$ has an integral basis of height $O(n^3(\lfloor 2log_q n\rfloor + \Delta))$. Now the bound stated for an arbitrary ideal over $\mathrm{Rad}(\mathcal{A})$ follows from the fact that any such ideal can be obtained as a sum of $\mathrm{Rad}(\mathcal{A})$ and some of the ideals I'_j. $\square$

We are now in a position to prove the central result of this section, stating that the Jacobson radical of finite dimensional associative algebras over finitely generated function fields of positive characteristic can be computed by a deterministic algorithm which is exponential only in the parameter m.

Theorem 3.4.5 *Let K be a finite extension of $\mathbb{F}_q(X_1, \ldots, X_m)$. Let $d := [K : \mathbb{F}_q(X_1, \ldots, X_m)]$. Suppose that $\mathcal{A}$, an n-dimensional algebra over K is given by structure constants of height not exceeding Δ. Suppose further that K is given by structure constants over $\mathbb{F}_q(X_1, \ldots, X_m)$ of height at most Γ. Then there is a deterministic algorithm that outputs a basis of the Jacobson radical of $\mathcal{A}$ at the cost of $((n + d + \Delta + \Gamma)^m + \log q)^{O(1)}$ bit operations.*

Proof We note that the radical (as the set of strongly nilpotent elements) does not change if we consider $\mathcal{A}$ as an algebra over a subfield K_0 of K such that $[K : K_0]$ is finite. Since K is a finite extension of $\mathbb{F}_q(X_1, \ldots, X_m)$, we may assume that $K = \mathbb{F}_q(X_1, \ldots, X_m)$. The height of the structure constants over the latter field is bounded by $2\Gamma + \Delta$. We use the algorithm outlined in Proposition 3.2. We may assume that the structure constants are integral. After each step, we can correct the basis of the ideal I_j to an integral basis of height $O((nd)^3(\lfloor 2log_q(nd) \rfloor + 2\Gamma + \Delta))$ This can be done by solving $(n + d + \Delta + \Gamma)^{O(m)}$ linear equations over the finite field $\mathbb{F}_q$. $\square$

Corollary 3.4.6 *In addition to the assumptions of Theorem 3.6 suppose that m is a constant (e.g., K is a global function field). Then $\mathrm{Rad}(\mathcal{A})$ can be computed in deterministic polynomial time.*

Proof In this case the bound of Theorem 3.4.5 is a polynomial of the input size. $\square$

Chapter 4

Wedderburn decomposition over $\mathbb{F}_q(X_1, \ldots, X_m)$

In this chapter — based on the paper [IRSz], joint work with Lajos Rónyai and Ágnes Szántó — we give an f-algorithm (a deterministic polynomial time reduction to factoring polynomials over the prime field $\mathbb{F}_p$) for decomposing a finite dimensional semisimple algebra over a finite extension K of the field $\mathbb{F}_q(X_1, \ldots, X_m)$. We actually reduce our problem to decomposing algebras over finite fields, and this problem can in turn be solved by the f-algorithm of Friedl and Rónyai [FR]. For large primes the algorithm presented here is an improved analogue of the method of Gianni, Miller and Trager [GMT], which was proposed for decomposition of algebras over $\mathbb{Q}$. The main idea is to find central idempotents in a factor modulo an appropriate maximal ideal I, lift them to idempotents modulo a sufficiently large power of I, and then find zero divisors of reasonable height. For small primes we present a deterministic polynomial time algorithm which is an analogue of Berlekamp's method (cf. [Ber3]) for factoring polynomials over finite fields.

An extremely simple standard idea is reducing the task of computing the simple components of a not necessarily commutative algebra to the decomposition of the center. It is clear, that the center of the algebra is a direct sum of the centers of its simple components and a component is an ideal generated by its center. The center of an algebra can be computed by solving a system of linear equations.

Let us recall some classical material related to idempotents (cf. [Ke], Sections 31 and 32). Recall that an element e in an arbitrary ring $\mathcal{A}$ is *idempotent* if $0 \neq e = e^2$. Two idempotents e, f in $\mathcal{A}$ are called *orthogonal* if $ef = fe = 0$. The sum of pairwise orthogonal idempotents is again an idempotent. An idempotent e in $\mathcal{A}$ is called *primitive* if there exists no idempotent $f \neq e$ with $ef = f$. This in turn is equivalent to that e is not a sum of two orthogonal idempotents. In the important case where $\mathcal{A}$ is an Artinian ring, every idempotent is a sum of finitely many orthogonal primitive idempotents.

Assume that $\mathcal{A}$ is a commutative Artinian ring. Then there are finitely many idempotents in $\mathcal{A}$, the primitive idempotents are pairwise orthogonal, and every idempotent can be uniquely written as a sum of primitive idempotents. Let J denote the Jacobson radical of $\mathcal{A}$. Then J is nilpotent. Let N be a positive integer such that $J^N = 0$. Two idempotents in $\mathcal{A}$ are equal or orthogonal iff modulo J they are equal or orthogonal, respectively. An idempotent e is primitive iff it is a primitive idempotent modulo J. Also, if $e_1 + J$ is an idempotent in $\mathcal{A}/J$ then there is a unique idempotent e in $e_1 + J$. The following iterative procedure (cf. [Re], Section 6c) is available for finding the lift e. Define the sequence $e_1, e_2, \ldots$ by the recurrence formula

$$e_{i+1} = 3e_i^2 - 2e_i^3.$$

Then $e = e_{\lceil \log_2 N \rceil}$.

Assume that $\mathcal{A}$ is a commutative semisimple finite dimensional K-algebra. Then the primitive idempotents of $\mathcal{A}$ are the identity elements of the minimal ideals of $\mathcal{A}$. Therefore (up to solving systems of linear equations over K) it is equivalent to find the primitive idempotents and the minimal ideals of $\mathcal{A}$. Also, if K_0 is any subfield of $\mathcal{A}$ then the K_0-subalgebra of $\mathcal{A}$ generated by the (primitive) idempotents is the sum of the K_0-subalgebras of the simple components of $\mathcal{A}$ generated by the identity elements of the simple components. If K is of positive characteristic p, then with $K_0 = \mathbb{F}_p$ this subalgebra is in fact the subring $\mathcal{B}$ generated by the (primitive) idempotents and can be characterized as

$$\mathcal{B} = \{x \in \mathcal{A} | x^p = x\}.$$

Since the map $x \mapsto x^p$ is $\mathbb{F}_p$-linear, this leads to an efficient way to compute $\mathcal{B}$ in the case when K is a finite field given in the dense representation over the prime field $\mathbb{F}_p$: One computes an $\mathbb{F}_p$-basis $a_1, \ldots, a_n$ of $\mathcal{A}$, then the elements $a_1^p - a_1, \ldots, a_n^p - a_n$ via fast exponentiation (based on repeated squaring), and finally $\mathcal{B}$ as the solution space (in the coordinates $x_1, \ldots, x_n$ w.r.t. the basis $a_1, \ldots, a_n$) of the system

$$\sum x_i(a_i^p - a_i) = 0$$

of linear equations over $\mathbb{F}_p$. Combining with the method of [FR], essentially this is Berlekamp's reduction of the problem of factoring polynomials over a finite field to factor polynomials that split into linear factors over $\mathbb{F}_p$.

If we attempt to generalize the method above to the case when $K = \mathbb{F}_q(X_1, \ldots, X_m)$, then, we need first to compute an $\mathbb{F}_p$-subspace $\mathcal{B}'$ of $\mathcal{A}$ of small dimension containing $\mathcal{B}$.

This can be done based on Lemma 4.1.2. Then we would need the powers a_i^p for an $\mathbb{F}_p$-basis $a_1, a_2, \ldots$ of $\mathcal{B}'$. Unfortunately, these elements can be efficiently computed only if p is small. If p is large, then the size of a_i^p (and the intermediate powers occurring in repeated squaring) may explode.

We treat the case of large characteristic in a different way. Lemma 4.1.2 makes possible to find a maximal ideal I of $\mathbb{F}_q[X_1, \ldots, X_m]$ such that we can efficiently compute modulo powers of I in an appropriate subring of $\mathcal{A}$ containing the idempotents. (In essence, we imitate computations in the maximal order of $\mathcal{A}$ over the localization $\mathbb{F}_q[X_1, \ldots, X_m]_I$.) For an appropriately chosen exponent h, we compute the primitive idempotents modulo I^h (from the primitive idempotents modulo I via the lifting method described above), and use them to find zero divisors in $\mathcal{A}$.

4.1 Algorithms

Let $\mathcal{A}$ be a semisimple algebra over the field $\mathbb{F}_q(X_1, \ldots, X_m)$. The results of Section 3.4 can be applied to obtain polynomial bounds on the height of the simple components of a semisimple algebra. Lemma 3.4.3 gives immediately the next statement.

Corollary 4.1.1 *Let $\mathcal{A}$ be an n-dimensional semisimple algebra over the field $K = \mathbb{F}_q(X_1, \ldots, X_m)$. Assume that the structure constants are from $\mathbb{F}_q[X_1, \ldots, X_m]$ and their height is limited by Δ. Then any ideal I of $\mathcal{A}$ has an integral basis of height at most* $2(n\lfloor 2log_q n\rfloor + (n+1)\Delta)$. □

If we use a successive decomposition of the algebra, then after computing an ideal $I \lhd \mathcal{A}$ we can compute a basis of I of height at most

$$2(n\lfloor 2log_q n\rfloor + (n+1)\Delta)$$

by solving a system of linear equations (of moderate size). This correction step keeps the heights always under the bound.

The task of computing the Wedderburn decomposition is equivalent (up to polynomial time reductions) to the problem of finding a zero divisor in a commutative semisimple algebra $\mathcal{A}$: If a is a zero divisor in $\mathcal{A}$, then $\mathcal{A}$ is a direct sum of the ideal generated by a and the annihilator

$$\mathrm{ann}(a) = \{x \in \mathcal{A} \mid xa = 0\}$$

of a.

As mentioned in the introductory part, our task is equivalent to decomposing the center of $\mathcal{A}$, therefore we may assume that $\mathcal{A}$ is a commutative semisimple algebra over $K = \mathbb{F}_q(X_1, \dots, X_m)$. We can also make another important simplification. It will be convenient to assume that the algebra is *separable* over K, i.e., the simple components of $\mathcal{A}$ are all finite separable extensions of K. This can be achieved as follows. Let $j := \lfloor \log_p n \rfloor$, where n is the dimension of $\mathcal{A}$ over K. Let us consider the Frobenius map defined on $\mathcal{A}$ by

$$\Phi_j \colon x \mapsto x^{p^j}.$$

It is clear, that Φ_j is a ring-endomorphism of $\mathcal{A}$ which extends the Frobenius endomorphisms of the simple components. It is also known (see [Bas], Propositions 2.4.9 and 2.5.13), that the image of each component will be contained in a separable extension of K. Thus, if we take the K-algebra generated by p^jth power of all elements of a basis of the algebra, then we obtain a separable subalgebra. On the other hand, the ideal structure is preserved: all the idempotents are left fixed. Since $p^j \leq n$, the p^jth powers of the basis elements are of polynomial size and can be computed in polynomial time and we can assume that our algebra is separable over K.

We use the *bilinear trace form* of the algebra (see also Chapter 6). It is a symmetric F-bilinear function defined by

$$\mathrm{tr}(x, y) := \mathrm{Tr}(xy)$$

where Tr is the trace of the regular representation. If $\{a_1, \dots, a_n\}$ is a basis of the algebra $\mathcal{A}$ over K, then the *discriminant* of the algebra corresponding to that basis is the determinant

$$\mathrm{disc}_{\{a_1,\dots,a_n\}}\mathcal{A} := \det\left(\mathrm{tr}(a_i, a_j)\right)_{i,j=1}^n.$$

It is well-known (see [Bas], pp. 166–168), that for any commutative separable algebra the bilinear trace form is *nondegenerate*, i.e., the discriminant is nonzero.

Since the height of the terms $\mathrm{Tr}(a_i a_j)$ is bounded by 2Δ, we have a bound $2n\Delta$ on the height of the discriminant, where Δ is a bound on the heights of the structure constants. This observation shows that the discriminant can be computed by Chinese Remaindering.

The following lemma gives an estimate of the height of idempotents in a commutative separable algebra. A common multiple of the denominators of the idempotents is also exhibited. The proof is independent from the proof of Corollary 4.1.1.

Lemma 4.1.2 *Let $\mathcal{A}$ be an n-dimensional commutative separable algebra over the field $K = \mathbb{F}_q(X_1, \dots, X_m)$. Assume that the structure constants with respect to the basis*

$\{a_1, \ldots, a_n\}$ are from $\mathbb{F}_q[X_1, \ldots, X_m]$ and their heights are limited by Δ. Then any idempotent $e \in \mathcal{A}$ lies in the $\mathbb{F}_q$-space

$$\{\sum_{i=1}^{n} \frac{\alpha_i}{D} a_i \mid \alpha_i \in \mathbb{F}_q[X_1, \ldots, X_m],\ \deg_{X_j} \alpha_i \le (3n-2)\Delta, \quad 1 \le j \le m\},$$

where D is the discriminant $\mathrm{disc}_{\{a_1,\ldots,a_n\}}\mathcal{A}$.

Proof Let e be a nontrivial idempotent in $\mathcal{A}$, $I = e\mathcal{A}$ the ideal generated by e and $J = \mathrm{ann}(e)$ be the annihilator of e. Then for an arbitrary $a \in \mathcal{A}$, the left (or, equivalently, the right) action of ea on I coincides with the action of a on I, while the action of ea on J is zero. The integrality assumption on the structure constants implies that the characteristic roots of a_j are integral elements over $\mathbb{F}_q[X_1, \ldots, X_m]$, hence the characteristic roots of ea_j are also integral over $\mathbb{F}_q[X_1, \ldots, X_m]$. This implies that $\mathrm{Tr}(ea_j) \in \mathbb{F}_q[X_1, \ldots, X_m]$ for $1 \le j \le n$ and the characteristic polynomial f of the action of ea_j on I divides the characteristic polynomial g of the action of a_j on $\mathcal{A}$. Let $\mathrm{Tr}_I(d)$ denote the trace of the action on I of an element $d \in I$. Clearly we have $\mathrm{Tr}_I(d) = \mathrm{Tr}(d)$ for every $d \in I$, in particular $\mathrm{Tr}_I(ea_j) = \mathrm{Tr}(ea_j)$ holds. We infer that $-\mathrm{Tr}(ea_j)$ is a coefficient of f, therefore the height of $\mathrm{Tr}(ea_j)$ is at most $n\Delta$, because this quantity is a bound on the height of the coefficients of g. On the other hand, $e = \sum_{i=1}^{n} \beta_i a_i$ is the unique solution of the following system of linear equations:

$$\sum_{i=1}^{n} \beta_i \mathrm{Tr}(a_i a_j) = \mathrm{Tr}(ea_j),\ j = 1, \ldots, n.$$

Cramer's rule implies the assertion. □

Let $\mathcal{M}$ denote the free $R = \mathbb{F}_q[X_1, \ldots, X_m]$-submodule of $\mathcal{A}$ with basis $\frac{1}{D}a_1, \ldots, \frac{1}{D}a_n$. We perform computations in $\mathcal{M}$ and work with the height of elements of $\mathcal{M}$ with respect to the R-basis above, i.e., for $u = \sum \alpha_i \frac{1}{D} a_i$,

$$\mathrm{height}(u) = \max\{\deg_{X_j} \alpha_i | i = 1, \ldots, n,\ j = 1, \ldots, m\}.$$

Obviously, for $u, v \in \mathcal{M}$, $\alpha \in \mathbb{F}_q[X_1, \ldots, X_m]$, we have

$$\mathrm{height}(\alpha u) \le \mathrm{height}(\alpha) + \mathrm{height}(u),$$

$$\mathrm{height}(u + v) \le \max\{\mathrm{height}(u), \mathrm{height}(v)\},$$

$$Duv \in \mathcal{M}, \quad \text{and} \quad \mathrm{height}(Duv) \le \Delta + \mathrm{height}(u) + \mathrm{height}(v).$$

We can settle now the case when p is small in comparison to the other parameters of the input. More precisely let us assume that $p \leq 2n\Delta$. In this case we can reduce the problem of finding the Wedderburn decomposition of $\mathcal{A}$ to the problem of factoring polynomials over finite fields in a way similar to Berlekamp's reduction [Ber3]. Let $\mathcal{B}$ denote the direct sum of the prime fields in the simple components of our (commutative separable) algebra $\mathcal{A}$. Clearly $\mathcal{B}$ is an $\mathbb{F}_p$-algebra consisting of the fixed points of the Frobenius endomorphism $x \mapsto x^p$ of the K-algebra $\mathcal{A}$. Moreover $\mathcal{B}$ is the $\mathbb{F}_p$-algebra generated by the idempotents of $\mathcal{A}$. From Lemma 4.1.2 it follows that every element $b \in \mathcal{B}$ can be written as a sum $b = \beta_1 a_1 + \cdots + \beta_n a_n$, where $\beta_i = \alpha_i / D$ for some $\alpha_i \in \mathbb{F}_q[X_1, \ldots, X_m]$ and $\deg_{X_j} \alpha_i \leq (3n-2)\Delta$. We infer that a basis over $\mathbb{F}_p$ of $\mathcal{B}$ can be obtained by solving a system of at most $(3pn\Delta)^m \log_p q$ linear equations in $(3n\Delta)^m \log_p q$ variables over $\mathbb{F}_p$. The coefficients of the equations can be obtained by fast exponentiation of the matrices of structure constants. Once we have a basis of $\mathcal{B}$, we can readily obtain structure constants over $\mathbb{F}_p$ and then compute the primitive idempotents of $\mathcal{B}$ (and therefore of $\mathcal{A}$) with the method of Friedl and Rónyai [FR, Ró2]. This completes the algorithm in the case $p \leq 2n\Delta$.

For a "large" prime p we follow the method of Gianni, Miller and Trager [GMT]. We shall, however, propose an improvement in the final step of constructing zero divisors from the lifted idempotents. As usual, for $f \in \mathbb{F}_q[X_1, \ldots, X_m]$ and $c_i \in \mathbb{F}_q$ we denote by $f(c_1, \ldots, c_m)$ the element of $\mathbb{F}_q$ obtained after substituting c_i for X_i.

Let us assume that $q > 2n\Delta$. We can find, using at most $m(2n\Delta + 1)$ tries, elements $c_1, \ldots, c_m \in \mathbb{F}_q$ such that $D(c_1, \ldots, c_m) \neq 0$. Indeed, we find first a substitution $X_1 = c_1$ from a set of $\mathbb{F}_q$ of cardinality $2n\Delta + 1$ such that $D(c_1, X_2, \ldots X_m) \not\equiv 0$. Then we repeat this for $X_2, \ldots, X_m$.

Let I be the ideal of R generated by the polynomials $X_1 - c_1, \ldots, X_m - c_m$. Then we have $R/I \cong \mathbb{F}_q$ and for every positive integer h, R/I^h is an F_q-algebra of dimension at most mh^m. (After a change of variables $Y_i = X_i - c_i$, we may assume that $c_i = 0$, for $i = 1, \ldots, m$. Then reducing a polynomial in R modulo I^h is equivalent to taking only the monomials of total degree less than h.) We also have that, for every h, the R/I^h-module $\mathcal{M}/I^h\mathcal{M}$ can be considered as an R/I^h algebra by taking multiplication modulo I^h. To be more precise, the structure constants with respect to the basis $\frac{1}{D}a_1, \ldots, \frac{1}{D}a_n$ will be the structure constants of the algebra $\mathcal{A}$ with respect to the basis $a_1, \ldots, a_n$ taken modulo I^h and multiplied by the inverse of D modulo I^h, which can be computed as follows: we write D as $D = c + f$ where $c \in \mathbb{F}_q$ and $f \in I$. Then $1 = (\frac{1}{c}D - \frac{1}{c}f)^h$ and hence we have to divide $(\frac{1}{c}D - \frac{1}{c}f)^h - (\frac{1}{c}f)^h$ by D. From this, we see that elements of $\mathcal{M}/I^h\mathcal{M}$ can be represented by $O(h^m n \log_q)$ bits and the R/I^h-algebra operations on $\mathcal{M}$ can be carried out in time $(h^m + n + \log_q)^{O(1)}$. In fact, $\mathcal{M}/I^h\mathcal{M}$ is an $\mathbb{F}_q$-algebra of dimension $n\dim_{\mathbb{F}_q} I^h$ with identity element $1_{\mathcal{A}} + I^h\mathcal{M}$.

If $0 \neq e$ is an idempotent of $\mathcal{A}$ then $e \in \mathcal{M}$ by Lemma 4.1.2. The definition of the multiplication gives immediately that $\bar{e}^2 = \bar{e}$, where $\bar{e}$ denotes the image of e in $\mathcal{M}/I^h\mathcal{M}$. Also, $e^2 = e \neq 0$ implies that $\bar{e} \neq 0$, for otherwise we would have the contradictory $e \in I^j\mathcal{M}$ for every $j \geq h$.

The case $h = 1$ deserves special attention. $\mathcal{M}/I\mathcal{M}$ is an n-dimensional $\mathbb{F}_q$-algebra. Since this algebra admits a nondegenerate bilinear trace form, it must be semisimple. (See [Di], and for more general results Chapter 3 in this dissertation.) It follows that for every integer $h \geq 1$, $I\mathcal{M} + I^h\mathcal{M}$ is the Jacobson radical of the finite ring $\mathcal{M}/I^h\mathcal{M}$. Using this fact, we can find the idempotents modulo I^h as follows. We compute the Wedderburn decomposition and hence the primitive idempotents of the $\mathbb{F}_q$-algebra $\mathcal{M}/I\mathcal{M}$ with the help of the f-algorithm of Friedl and Rónyai [FR, Ró2]. To be more specific, we obtain elements $f_1, \ldots, f_r \in \mathcal{M}$, such that $f_1 + I\mathcal{M}, \ldots, f_r + I\mathcal{M}$ are the primitive idempotents of the ring $\mathcal{M}/I\mathcal{M}$. Then we can lift these primitive idempotents to primitive idempotents of $\mathcal{M}/I^h\mathcal{M}$ by the classical method described in the introductory part in time $((\text{height}\mathcal{A}+h)^m \log q)^{O(1)}$. We obtain elements $e_1, \ldots, e_r \in \mathcal{M}$, that are reduced modulo I^h, i.e., with coordinates of total degree less than h, such that $e_1 + I^h\mathcal{M}, \ldots, e_r + I^h\mathcal{M}$ are the primitive idempotents in the ring $\mathcal{M}/I^h\mathcal{M}$.

An idempotent of $\mathcal{A}$ in $\mathcal{M}$ is also idempotent modulo I^h, whence it is a sum of some primitive idempotents of $\mathcal{M}/I^h\mathcal{M}$. If there are no nontrivial idempotents modulo I^h, i.e., $\mathcal{M}/I\mathcal{M}$ is simple, we can stop, our algebra is simple. On the other hand, if $h \geq (3n-2)\Delta$, then every idempotent $e \in \mathcal{M}$ can be obtained as a sum of a subset of $\{e_1, \ldots, e_r\}$.

Unfortunately, it can well happen that there are more idempotents in $\mathcal{M}/I^h\mathcal{M}$ than in $\mathcal{A}$ and then we should have to examine exponentially many sums as in [GMT]. The trick below helps. Note that an analogous method (using LLL basis reduction) could also be applied in the situation of [GMT].

If there are nontrivial idempotents in $\mathcal{M}$, then there is one, say e, which is a sum of a subset of $e_2, \ldots, e_n$. From Lemma 4.1.2 we infer that the $\mathbb{F}_q$-subspace $\mathcal{M}_1$ of $\mathcal{M}$ generated by $e_2, \ldots, e_n$ contains an element of height at most $(3n-2)\Delta$. On the other hand, we claim that if h is large enough (actually, $h > m(3n^2 - n - 1)\Delta$), then a nonzero element of $\mathcal{M}_1$ of height at most $(3n-2)\Delta$ is in fact a zero divisor in $\mathcal{A}$. Indeed, assume that $h > m(3n^2 - n - 1)\Delta$ and let

$$0 \neq u = \sum_{i=2}^{r} \gamma_i e_i \in \mathcal{M}_1$$

with $\gamma_2 \ldots, \gamma_r \in \mathbb{F}_q$ and $\text{height}(u) \leq (3n-2)\Delta$. Let $\bar{u}$ denote the residue class $u + I^h\mathcal{M}$ of u modulo $I^h\mathcal{M}$. Let $\beta_1, \ldots, \beta_t$ be an enumeration of the distinct nonzero elements from

$\gamma_2, \ldots, \gamma_r$. It is straightforward to see that in the $\mathbb{F}_q$-algebra $\mathcal{M}/I^h\mathcal{M}$,

$$\bar{u}\prod_{i=1}^{t}(\bar{u}-\beta_i 1)=0, \quad \text{but} \quad \prod_{i=1}^{t}(\bar{u}-\beta_i 1)\neq 0.$$

Using the definition of multiplication in $\mathcal{M}/I^h\mathcal{M}$ we obtain

$$D^t u\prod_{i=1}^{t}(u-1_{\mathcal{A}})\in I^h\mathcal{M}, \quad \text{but} \quad \prod_{i=1}^{t}(u-1_{\mathcal{A}})\neq 0.$$

Since $1_{\mathcal{A}}$ is also an idempotent, we have $\text{height}(u-\beta_i 1_{\mathcal{A}})\leq(3n-2)\Delta$, whence

$$\text{height}(D^t u\prod_{i=1}^{t}(u-\beta_i 1_{\mathcal{A}}))\leq t\Delta+\text{height}(u)+\sum_{i=1}^{t}\text{height}(u-\beta_i 1_{\mathcal{A}})\leq$$

$$t\Delta+(t+1)(3n-2)\Delta\leq(3n^2-n-1)\Delta.$$

From the fact that $I^h\mathcal{M}$ does not contain elements of total degree less than $h>m(3n^2-n-1)\Delta$, we infer

$$u\prod_{i=1}^{t}(u-1_{\mathcal{A}})=0,$$

i.e., u is a zero divisor in $\mathcal{A}$.

As a conclusion of the preceding argument, if we can find an element $u\neq 0$ of total degree not greater than $m(3n-2)\Delta$ in the $\mathbb{F}_q$-space $\mathcal{M}_1$ generated by the reduced lifted idempotents $e_2, \ldots, e_n$, then we have a zero divisor in the algebra $\mathcal{A}$. An element u with these properties can be found (if exists) by solving a system of linear equation over $\mathbb{F}_q$. If there is no nonzero solution, then we conclude that $\mathcal{A}$ is simple. If we find a solution u, then we can split the algebra to the proper ideal $\mathcal{A}u$ and to its complement.

An iteration of this method gives the minimal ideals of $\mathcal{A}$ in at most n rounds. Note that we do not need to recompute the lifted idempotents for the ideals. This finishes the description of our f-algorithm for the case $K=\mathbb{F}_q(X_1, \ldots, X_m)$. The overall cost of the method is $((n+\Delta)^m+\log q)^{O(1)}$ bit operations. We can easily extend this result to the case when K is a finite extension of $\mathbb{F}_q(X_1, \ldots, X_m)$. We consider $\mathcal{A}$ as an algebra over $\mathbb{F}_q(X_1, \ldots, X_m)$. This change does not affect the set of the central primitive idempotents. Also, structure constants over $\mathcal{A}$ can be computed efficiently. We have the following.

Theorem 4.1.3 *Assume that $\mathcal{A}$ is an n-dimensional semisimple algebra over a finite extension K of the function field $\mathbb{F}_q(X_1, \ldots, X_m)$ of degree $d=[K:\mathbb{F}_q(X_1, \ldots, X_m)]$. The algebra is given by structure constants of height (maximum degree in each variable) limited* *by Δ. Suppose further that K is given by structure constants over $\mathbb{F}_q(X_1, \ldots, X_m)$ of height at most Γ. Then there is an f-algorithm that computes (bases of) the minimal ideals of $\mathcal{A}$ using $((n+d+\Delta+\Gamma)^m + \log q)^{O(1)}$ bit operations.* $\square$

Using the fact that we can factor a polynomial $f(X) \in \mathbb{F}_p[X]$ in randomized polynomial time (cf. Berlekamp [Ber3]), we obtain the following result.

Corollary 4.1.4 *Under the assumptions of Theorem 4.1.3, the Wedderburn decomposition of $\mathcal{A}$ can be computed in Las Vegas time $((n+d+\Delta+\Gamma)^m + \log q)^{O(1)}$.* $\square$

Corollary 4.1.5 *In addition to the assumptions of Theorem 4.1.3 suppose that m is a constant, e.g., K is a global function field. Then the Wedderburn decomposition of $\mathcal{A}$ can be computed by a polynomial time f-algorithm or a polynomial time Las Vegas algorithm.* $\square$

We can make a further improvement to obtain a deterministic algorithm for computing $\mathcal{B}$, the subring generated by the idempotents of $\mathcal{A}$.

We can compute the $\mathbb{F}_p$-subspace $\mathcal{L}$ of $\mathcal{M}$ generated by $e_1, \ldots, e_r$ without knowing $e_1, \ldots, e_r$ explicitly as follows. Since R/I^h is a finite dimensional $\mathbb{F}_p$-algebra, the R/I^h-algebra $\mathcal{M}/I^h\mathcal{M}$ is a finite dimensional $\mathbb{F}_p$-algebra as well. Let $\bar{\mathcal{L}}$ denote the image of $\mathcal{L}$ at the natural homomorphism $\mathcal{M} \to \mathcal{M}/I^h\mathcal{M}$. Observe that

$$\bar{\mathcal{L}} = \{x \in \mathcal{M}/I^h\mathcal{M} \,|x^p = x\}.$$

First we compute $d_1, \ldots, d_r \in \mathcal{M}$, such that $d_1 + I\mathcal{M}, \ldots, d_r + I\mathcal{M}$ is a basis of $\{x \in \mathcal{M}/I\mathcal{M} \,|x^p = x\}$. This can be done in polynomial time via solving a system of linear equations over $\mathbb{F}_p$. From the fact $\mathrm{Rad}(\mathcal{M}/I^h\mathcal{M}) = I\mathcal{M} + I^h\mathcal{M}$ we see that with $l = \lceil \log_p h \rceil$, the system $d_1^{p^l} + I^h\mathcal{M}, \ldots, d_r^{p^l} + I^h\mathcal{M}$ forms an $\mathbb{F}_p$-basis of $\bar{\mathcal{L}}$. We can compute $d_i^{p^l}$ modulo I^h using fast exponentiation in $\mathcal{M}/I^h\mathcal{M}$. This gives a system $b_1, \ldots, b_r$ of r elements of $\mathcal{M}$ with coordinates of total degree less than h, that is an $\mathbb{F}_p$ basis $\mathcal{L}$. We prove that, if h is large enough, then $\mathcal{B}$ coincides with the $\mathbb{F}_p$-subspace of $\mathcal{L}$ consisting of the elements of "small" height.

We set

$$\mathcal{B}' = \{u \in \mathcal{L} | \mathrm{height}(u) \leq (3n-2)\Delta\}.$$

Obviously, $\mathcal{B}'$ can computed from $\mathcal{L}$ in time $(h^m + n + \log q)^{O(1)}$ as the solution space of a system of homogeneous linear equations over $\mathbb{F}_p$. By Lemma 4.1.2 for every h we have $\mathcal{B} \subseteq \mathcal{B}'$. We show that under the assumption $h > m(3n^2 - n - 1)\Delta$, $\mathcal{B}' \subseteq \mathcal{B}$ also holds. Let

$u \in \mathcal{B}'$. Then there exist $\beta_1, \ldots, \beta_r \in \mathbb{F}_p$ such that $u = \sum \beta_i e_i$. It is straightforward to see that with $\bar{u} = u + I^h\mathcal{M}$ we have $\prod_{i=1}^r (\bar{u} - \beta_i 1) = 0$ in the ring $\mathcal{M}/I^h\mathcal{M}$, whence

$$D^{r-1} \prod_{i=1}^{r} (u - \beta_i 1_{\mathcal{A}}) \in I^h\mathcal{M}.$$

Observe that for every $i = 1, \ldots, r$, (since $1_{\mathcal{A}} \in \mathcal{B} \subseteq \mathcal{B}'$) we have $u - \beta_i 1_{\mathcal{A}} \in \mathcal{B}'$, whence

$$\text{height}(D^{r-1} \prod_{i=1}^{r} (u - \beta_i 1_{\mathcal{A}})) \leq (r-1)\Delta + \sum_{i=1}^{r} \text{height}(u - \beta_i 1_{\mathcal{A}}) \leq$$

$$(r-1)\Delta + r(3n-2)\Delta \leq (3n^2 - n - 1)\Delta.$$

Since $I^h\mathcal{M}$ does not contain nonzero elements of total degree less than h, we have

$$\prod_{i=1}^{r} (u - \beta_i 1_{\mathcal{A}}) = 0.$$

But then the roots of the minimal polynomial of u are all in $\mathbb{F}_p$, whence $u \in \mathcal{B}$.

We have proved the following.

Theorem 4.1.6 *Assume that $\mathcal{A}$ is an n-dimensional semisimple algebra over a finite extension K of the function field $\mathbb{F}_q(X_1, \ldots, X_m)$ of degree $d = [K : \mathbb{F}_q(X_1, \ldots, X_m)]$ given by structure constants of height (maximum degree in each variable) limited by Δ. Suppose further that K is given by structure constants over $\mathbb{F}_q(X_1, \ldots, X_m)$ of height at most Γ. Then there is a deterministic algorithm running in time $((n + d + \Delta + \Gamma)^m + \log q)^{O(1)}$ that computes the $\mathbb{F}_p$-subalgebra $\mathcal{B}$ generated by the central idempotents of $\mathcal{A}$.* $\square$

Chapter 5

Cartan subalgebras

This chapter is based on the paper [GIR], joint work with Willem A. de Graaf and Lajos Rónyai. We consider the algorithmic problem of computing Cartan subalgebras in Lie algebras over finite fields and global fields. We present a deterministic polynomial time algorithm for the case when the ground field K is sufficiently large. Our method is based on a solution of a linear algebra problem: the task of finding a locally regular element in a subspace of linear transformations. Also, we give a polynomial time algorithm for restricted Lie algebras over arbitrary finite fields. Both methods require an auxiliary procedure for finding non-nilpotent elements in subalgebras. This problem is also treated.

Throughout the chapter K denotes a field and $\mathcal{L}$ a finite dimensional Lie algebra over K; we write $n = \dim_K \mathcal{L}$. For the basic definitions and results on Lie algebras the reader is referred to [Hum] and [Jac]. A subalgebra $\mathcal{H} \leq \mathcal{L}$ is a *Cartan subalgebra* of $\mathcal{L}$ if $\mathcal{H}$ is nilpotent and equals its normalizer: $N_{\mathcal{L}}(\mathcal{H}) = \mathcal{H}$. Via decompositions into root spaces, Cartan subalgebras proved to be extremely useful in exploring the structure of Lie algebras.

The main result of this chapter (in Section 5.4) is a deterministic polynomial time algorithm for computing Cartan subalgebras in the following two cases:

1. K is a global field or a finite field with $|K| > \dim_K \mathcal{L}$.

2. K is an arbitrary finite field and $\mathcal{L}$ is a restricted Lie algebra.

We point out, that if $|K| \gg \dim_K \mathcal{L}$, then a very efficient randomized method is also available.

In the case $|K| > \dim_K \mathcal{L}$ we consider a much more general problem, which, we believe, is interesting on its own right. Let V denote a linear space over the field K, $\dim_K V = n$. Let $\mathrm{End}_K V$ denote the set of K-linear transformations (K-endomorphisms) of V. For an endomorphism $a \in \mathrm{End}_K V$ we denote by $V_0(a)$ the Fitting null component of a on V:

$$V_0(a) := \{v \in V;\quad a^m v = 0 \text{ for some positive integer } m\}.$$

It is immediate that $V_0(a)$ is a subspace of V, in fact, it is the largest subspace of V on which a acts as a nilpotent endomorphism.

Let $U \leq V$ be a subspace of V and $D \leq \mathrm{End}_K V$ be a K-subspace of $\mathrm{End}_K V$. We denote by $N_D(U)$ the subspace of D leaving U invariant:

$$N_D(U) := \{a \in D |\ aU \subseteq U\}.$$

An element $a \in D$ is *locally regular* (in D) if every element of $N_D(V_0(a))$ acts as a nilpotent endomorphism on $V_0(a)$. We give a deterministic procedure for finding a locally regular element in a given $D \leq \mathrm{End}_K V$, provided that $|K| > \dim_K V$ (Section 5.3). The procedure runs in polynomial time, if we can efficiently find non-nilpotent transformations in subspaces of form $N_D(V_0(a))$, where $a \in D$. The algorithm for locally regular elements is applicable to Lie algebras via the adjoint representation. It turns out that $V_0(\mathrm{ad}_{\mathcal{L}} a)$ is a Cartan subalgebra of $\mathcal{L}$, if $\mathrm{ad}_{\mathcal{L}} a$ is a locally regular element of $\mathrm{ad}(\mathcal{L})$.

Our approach requires a solution of an important subtask: for a given subalgebra $N \leq \mathcal{L}$ find an element $a \in N$ such that $\mathrm{ad}_N a$ is not a nilpotent map. Polynomial time algorithms for this problem are considered in Section 5.2.

If K is infinite, then, in addition to the number of arithmetical operations, we have to bound the size of the numbers we work with. This will be done with the aid of the Reduction Lemma (Section 5.1), which allows us to keep the coefficients of the elements small during the computation.

In Section 5.5 we present applications to semisimple associative algebras. In the literature several polynomial time randomized algorithms working in semisimple algebras are known where randomization occurs only at the point of finding so-called *splitting elements*. Roughly speaking, the minimal polynomial of a splitting element has as many distinct roots (in the algebraic closure of the ground field) as possible. Our methods give a deterministic polynomial time algorithm for finding splitting elements in separable associative algebras over finite fields or global fields, generalizing a result of Rónyai [Ró5].

5.1 The Sparse Zeros Lemma and a reduction procedure

In symbolic computation it is a frequent task to find elements $\alpha_1, \alpha_2, \ldots \alpha_n \in K$ efficiently such that $f(\alpha_1, \ldots \alpha_n) \neq 0$, where $f \in K[X_1, \ldots, X_n]$ is a nonzero polynomial. The following statement by J. T. Schwartz [Sch] and R. E. Zippel [Z] provides a powerful randomized solution.

Sparse Zeros Lemma. *Let $f(X_1, X_2, \ldots, X_n) \in K[X_1, X_2, \ldots, X_n]$ be a polynomial, $\deg f = d$. Let Ω be a subset of K, $|\Omega| = N$. Then the number of vectors $u = (\alpha_1, \ldots, \alpha_n) \in \Omega^n$ for which $f(u) = 0$ is at most dN^{n-1}.*

Thus, if we select a random vector $u \in \Omega^n$ with uniform distribution, then the probability of $f(u)$ being 0 is at most d/N. By working with a sufficiently large set Ω, we can make this probability arbitrarily small.

The Sparse Zeros lemma implies that if $N > d$, then there is an $u \in \Omega^n$ for which $f(u) \neq 0$. To our knowledge, there is no efficient deterministic algorithm for finding such vectors in general. We have however the following.

Lemma 5.1.1 (Reduction Lemma) *Suppose that we have a polynomial $f \in K[X_1, X_2, \ldots, X_n]$, $\deg f = d$ and a subset $\Omega \subset K$ with $d < |\Omega| = N$. Suppose further that we are given a substitution $(\beta_1, \beta_2, \ldots, \beta_n) \in K^n$ with $f(\beta_1, \beta_2, \ldots, \beta_n) \neq 0$. Then we can find a $v \in \Omega^n$ for which $f(v) \neq 0$, at the expense of at most $n(d+1)$ tests of the form $f(u) \stackrel{?}{=} 0$ on vectors $u \in (\Omega \cup \{\beta_1, \beta_2, \ldots, \beta_n\})^n$.*

Proof Starting with $v^{(0)} = (\beta_1, \beta_2, \ldots, \beta_n)$, we construct a sequence of vectors $v^{(1)}, \ldots, v^{(n)}$, where $v^{(i)}$ is of the form $v^{(i)} = (\alpha_1, \ldots \alpha_i, \beta_{i+1}, \ldots, \beta_n)$, $\alpha_1, \ldots, \alpha_n \in \Omega$, and $f(v^{(i)}) \neq 0$ for $0 \leq i \leq n$. In particular $v = v^{(n)}$ will have the required properties. We describe how to obtain $v^{(i)}$ from $v^{(i-1)}$. Consider the univariate polynomial $g_i(X) = f(\alpha_1, \ldots \alpha_{i-1}, X, \beta_{i+1}, \ldots, \beta_n) \in K[X]$. We have $\deg g_i \leq d$ and g_i is not identically zero because $g_i(\beta_i) = f(v^{(i-1)}) \neq 0$ by assumption. These imply that among arbitrarily selected $d+1$ elements of Ω there must be one, which we denote by α_i, for which $g_i(\alpha_i) \neq 0$. We can therefore construct $v^{(i)}$ from $v^{(i-1)}$ with at most $d+1$ tests and the statement follows. $\square$

The Reduction Lemma is a tool to control the sizes of coefficients in the algorithms. (In subsequent applications to cases $\mathrm{char}K = 0$, Ω will be a set of small integers. Similarly, for a global function field K, Ω will be a set of polynomials of small degree.)

Suppose that we are given a subspace $D \leq \mathrm{End}_K V$ of linear transformations of V by a basis $a_1, \ldots, a_l$. Suppose further that we have an element $a = \gamma_1 a_1 + \cdots + \gamma_l a_l \in D$ such that $\dim_K V_0(a) = r$. Let Ω be a subset of K, $|\Omega| = n+1$, where $n = \dim_K V$. With the aid of the Reduction Lemma we can construct an element $b = \alpha_1 a_1 + \cdots + \alpha_l a_l$ with $\alpha_i \in \Omega$ and $\dim_K V_0(b) \leq r$. To see this, consider a "generic element" $A = X_1 a_1 + X_2 a_2 + \cdots + X_l a_l$, where the X_j are indeterminates over K. Let $f(Y) \in K(X_1, \ldots, X_l)[Y]$ stand for the characteristic polynomial of the linear transformation A of $V \otimes_K K(X_1, \ldots, X_l)$ as a linear space over $K(X_1, \ldots, X_l)$. We have

$$f(Y) = \det(YI_n - A) = Y^n + f_1 Y^{n-1} + \cdots + f_{n-1} Y + f_n,$$

where $f_i \in K[X_1, X_2, \ldots, X_l]$. Moreover f_i is a homogeneous polynomial and if $f_i \neq 0$, then $\deg f_i = i$.

It is also clear that the characteristic polynomial of an element $c = \delta_1 a_1 + \cdots + \delta_l a_l \in D$ over K is obtained by making the substitutions $x_i = \delta_i$. The statement below is immediate.

Proposition 5.1.2 *Let $c = \delta_1 a_1 + \cdots + \delta_l a_l$ be an element of $\mathrm{End}_K V$ with characteristic polynomial $h(y) \in K[y]$. The following are equivalent:*
(a) $\dim_K V_0(c) = s$.
(b) s is the largest integer such that y^s divides $h(y)$.
(c) $f_{n-s}(\delta_1, \delta_2, \ldots, \delta_l) \neq 0$, and $f_j(\delta_1, \delta_2, \ldots, \delta_l) = 0$ for $n - s < j \leq n$.
(d) $s = n - \mathit{rank}(c^n)$. □

We can apply the Reduction Lemma. With at most $(n - r + 1)l \leq (n + 1)n$ tests involving values from the set $\Omega \cup \{\gamma_1, \ldots, \gamma_l\}$ we find an element $b = \alpha_1 a_1 + \cdots + \alpha_l a_l$ such that $\alpha_i \in \Omega$ and $f_{n-r}(\alpha_1, \alpha_2, \ldots, \alpha_l) \neq 0$, and hence $\dim_K V_0(b) \leq r$. A test can be done by computing and inspecting the characteristic polynomial of the endomorphism in question. Alternatively, by Proposition 5.1.2 (d) one can test an element c by computing the rank of c^n. We have the following.

Corollary 5.1.3 *Let K be a global field. In case K is a number field, we set $\Omega = \{1, 2, \ldots, n + 1\}$. In case K is a finite extension of $\mathbb{F}_q(X)$, we assume that Ω consists of polynomials of degree at most $\lceil \log_q(n + 2) \rceil$. Suppose that we are given a subspace $D \leq \mathrm{End}_K V$ of linear transformations of V ($\dim_K V = n$) by a basis $a_1, \ldots, a_l$. Suppose also that we have an element $a = \gamma_1 a_1 + \cdots + \gamma_l a_l \in D$ such that $\dim_K V_0(a) = r$. Then we can find an element $b = \alpha_1 a_1 + \cdots + \alpha_l a_l$ with $\alpha_i \in \Omega$ and $\dim_K V_0(b) \leq r$ in deterministic polynomial time.*

Proof It suffices to observe that the tests required by the preceding method can all be performed in time polynomial in the sizes of the matrices a_i and the coefficients γ_j. □

5.2 Non-nilpotent elements in Lie algebras

Suppose that we are given a Lie algebra $\mathcal{L}$ by structure constants over the field K. Let $a_1, a_2, \ldots a_n$ be the input basis of $\mathcal{L}$. Our objective is to find an element $a \in \mathcal{L}$ such that $\mathrm{ad}_{\mathcal{L}} a$ is not a nilpotent endomorphism of $\mathcal{L}$, provided that $\mathcal{L}$ is not a nilpotent algebra.

The first method we describe is an iterative algorithm. Suppose that we have a subset $B_r = \{b_1, b_2, \ldots, b_r\}$, $(r < n)$, of linearly independent elements of $\mathcal{L}$ such that the linear space KB_r generated by B_r over K is a (Lie) subalgebra of $\mathcal{L}$ which acts nilpotently on $\mathcal{L}$. The latter condition means that $\mathrm{ad}_{\mathcal{L}}c_1\mathrm{ad}_{\mathcal{L}}c_2\cdots\mathrm{ad}_{\mathcal{L}}c_n = 0$ (as a product of linear transformations of the K-space $\mathcal{L}$) whenever $c_i \in KB_r$. The following sub-algorithm either produces a larger such set B_{r+1}, or finds a non-nilpotent element of $\mathcal{L}$.

1. Find an element $b \in \mathcal{L} \setminus KB_r$ such that $[b, B_r] \subseteq KB_r$.
This can be done either by solving a system of linear equations, or by letting b be first an element of the input basis which is in $\mathcal{L} \setminus KB_r$ and then repeatedly replacing b by $[b, b_i]$ if $[b, b_i] \notin KB_r$. As KB_r acts nilpotently on $\mathcal{L}$, we need no more than $n-1$ such replacement steps.

2. Check if $\mathrm{ad}_{\mathcal{L}}b$ is a nilpotent endomorphism of $\mathcal{L}$. If not, then output b as a non-nilpotent element, otherwise set $b_{r+1} = b$ and $B_{r+1} = \{b_1, b_2, \ldots, b_{r+1}\}$.

We show that if b is ad-nilpotent upon termination, then B_{r+1} is again an independent set and KB_{r+1} is a subalgebra acting nilpotently on $\mathcal{L}$. The first claim is immediate, as $b \notin KB_r$ by construction. As for the second, we have $[B_{r+1}, B_{r+1}] \subseteq [b, B_r] \cup [B_r, b] \cup [B_r, B_r] \subseteq KB_r$, giving that B_{r+1} is a subalgebra. In fact, we have obtained the slightly stronger fact that the subset $U = \mathrm{ad}_{\mathcal{L}}(\{b\} \cup KB_r) \leq \mathrm{End}_K\mathcal{L}$ is closed under the bracket operation. Moreover the elements of U are all nilpotent maps of $\mathcal{L}$. By [Jac], p. 33, Theorem 1, the associative algebra $\mathcal{A}$ generated by U is nilpotent, hence $\mathrm{ad}_{\mathcal{L}}(KB_{r+1}) \leq \mathcal{A}$ is also nilpotent.

Our method for finding a non-nilpotent element starts with the set $B_0 = \emptyset$ and repeatedly applies the sub-algorithm until either a non-nilpotent element is found, or $r = n$ is attained. In this case $KB_n = \mathcal{L}$ is a nilpotent algebra and therefore every element of $\mathcal{L}$ is ad-nilpotent.

The number of arithmetical operations is bounded by a polynomial of n, hence it gives a polynomial time algorithm if K is finite. If K is infinite then we have no satisfactory bound on the size (of the coefficients) of the elements b_i. Nevertheless, in the practically very important case $\mathrm{char}K = 0$ a quite simple method is available.

Proposition 5.2.1 *Let $\mathcal{L}$ be a Lie algebra over a field K of characteristic zero. Let $a_1, a_2, \ldots a_n$ be a basis of $\mathcal{L}$ over K. If $\mathcal{L}$ is not a nilpotent algebra, then the set $\{a_1, \ldots, a_n\} \cup \{a_i + a_j;\ 1 \leq i < j \leq n\}$ contains an element which is not ad-nilpotent.*

Proof If $\mathcal{L}$ is solvable but not nilpotent, then by [Jac], p. 45, Corollary 2, one of the basis elements a_i is not ad-nilpotent. If $\mathcal{L}$ is not solvable then [Jac], p. 73, Theorem 5 implies that the Killing form of $\mathcal{L}$ is not identically zero: there exist basis elements a_i, a_j $(i \leq j)$

such that $\mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_i \mathrm{ad}_{\mathcal{L}} a_j) \neq 0$. From

$$\mathrm{Tr}((\mathrm{ad}_{\mathcal{L}} a_i + \mathrm{ad}_{\mathcal{L}} a_j)^2) - \mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_i^2) - \mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_j^2) =$$

$$= \mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_i \mathrm{ad}_{\mathcal{L}} a_j) + \mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_j \mathrm{ad}_{\mathcal{L}} a_i) = 2\mathrm{Tr}(\mathrm{ad}_{\mathcal{L}} a_i \mathrm{ad}_{\mathcal{L}} a_j) \neq 0$$

we infer that the elements $a_i + a_j$, a_i, a_j cannot be all ad-nilpotent. □

Proposition 5.2.1 provides an efficient method for finding a non-nilpotent element if K is a number field. We can solve the problem by inspecting at most $\binom{n+1}{2}$ elements of $\mathcal{L}$. These elements have small coefficients. If $\mathcal{L}$ is semisimple then we can do better. Using the fact that the Killing form is non-degenerate on semisimple algebras, the argument of Proposition 5.2.1 gives that the set $\{a_1, a_2, \ldots, a_n, a_1 + a_2, \ldots, a_1 + a_n\}$ contains a non-nilpotent element.

Assume now that $K = \mathbb{F}_q(X)$ and $a_1, \ldots, a_n$ is a basis of $\mathcal{L}$ such that the structure constants w.r.t. this basis are in $\mathbb{F}_q[X]$. Such a basis can be obtained using the standard trick of clearing the denominators. Assume further that the degrees of structure constants are at most d. Let Λ be the $\mathbb{F}_q[X]$-submodule of $\mathcal{L}$ generated by $a_1, \ldots, a_n$. Then Λ is a free $\mathbb{F}_q[X]$-module with basis $a_1, \ldots, a_n$ and it is in fact a Lie subalgebra over the ring $\mathbb{F}_q[X]$ of $\mathcal{L}$ such that $\mathbb{F}_q(X)\Lambda = \mathcal{L}$. Obviously, $\mathcal{L}$ is nilpotent iff $\Lambda^{n+1} = 0$. $\mathcal{L}^{n+1}$ has a basis consisting of elements of the form $[b_1, [b_2, \ldots [b_n, b_{n+1}] \ldots]]$, where $b_1, \ldots, b_n, b_{n+1} \in \{a_1, \ldots, a_n\}$. The coefficients of these elements, expressed w.r.t. the basis $a_1, \ldots, a_n$, are polynomials of degree at most nd. Let $f(X) \in \mathbb{F}_q[X]$ be an arbitrary polynomial. The factor module $\mathcal{L}_f = \Lambda / f(X)\Lambda$ inherits the Lie-ring structure of Λ. $\mathcal{L}_f$ is in fact a Lie algebra of rank n over the factor ring $\mathbb{F}_q[X]/(f(X))$ (the structure constants are reduced modulo f). Since $\mathbb{F}_q[X]/(f(X))$ is an associative F_q-algebra, we have that $\mathcal{L}_f$ is a Lie algebra over F_q of dimension $n\deg f$ and an $\mathbb{F}_q$-basis of $\mathcal{L}_f$ together with the structure constants can be computed efficiently. Obviously, if $\mathcal{L}_f$ turns out to be non-nilpotent and we find an element $\bar{a} \in \mathcal{L}_f$ which is not ad-nilpotent, then any lift a of $\bar{a}$ to Λ is not ad-nilpotent either. On the other hand, we have that if $\deg f > nd$ then $\mathcal{L}$ is nilpotent iff $\mathcal{L}_f$ is nilpotent. Therefore if we choose an arbitrary polynomial $f(X) \in \mathbb{F}_q[X]$ of degree $nd + 1$, then, using the iterative method for $\mathcal{L}_f$, we can decide in deterministic polynomial time whether $\mathcal{L}$ is nilpotent. Moreover, if $\mathcal{L}$ is not nilpotent, then find an element $a \in \Lambda$ such that $\bar{a} = a + f(X)\Lambda$ is a non-nilpotent element of $\mathcal{L}_f$.

An alternative (and probably more efficient) method is to test successively $\mathcal{L}_{f_1}, \mathcal{L}_{f_2}, \ldots$, where $f_1, f_2, \ldots$ are distinct irreducible polynomials in $\mathbb{F}_q[X]$. The procedure either stops with a non-nilpotent element modulo f_i, or with the conclusion that $\mathcal{L}$ is nilpotent if $\deg f_1 + \deg f_2 + \ldots + \deg f_i > nd$ and $\mathcal{L}_{f_1}, \mathcal{L}_{f_2}, \ldots, \mathcal{L}_{f_i}$ are nilpotent.

We summarize the results of this section in the following statement.

Theorem 5.2.2 *Let K be either a finite field or a global field. Suppose that we are given a non-nilpotent Lie algebra by structure constants over K. Then we can find in deterministic polynomial time an element $a \in \mathcal{L}$ such that $\mathrm{ad}_{\mathcal{L}} a$ is not a nilpotent linear transformation of $\mathcal{L}$.* $\square$

Remarks 1. If $\mathcal{L}$ is not nilpotent and $|K| \geq cn$ for a constant $c > 1$, then if we select coefficients α_i uniformly and independently from a subset $\Omega \subset K$, $|\Omega| \geq cn$, then the element $a = \alpha_1 a_1 + \cdots + \alpha_n a_n$ will be non-nilpotent with probability at least $1 - 1/c$. This follows readily from the Sparse Zeros Lemma and Proposition 5.1.2. Thus, if K is sufficiently large, then we have a good randomized method.
2. Our methods are applicable to subalgebras as well, where the input is represented in a different way: instead of structure constants we have basis vectors from a larger structure.

5.3 Locally regular endomorphisms

In this section we consider a linear algebra problem which includes a major subcase of the task of computing Cartan subalgebras. We believe that the problem is interesting on its own right, therefore we give here a separate treatment. Our method is an extended algorithmic version of the existence argument for Cartan subalgebras given on pp. 78–80, by Humphreys [Hum].

Suppose that we are given a subspace $D \leq \mathrm{End}_K V$ ($\dim_K V = n$) by a basis $a_1, a_2, \ldots a_l$. We wish to find a locally regular element $a \in D$.

We assume the existence of an auxiliary procedure *NON-NILP*(a), which, on input $a \in D$ returns an element $b \in N_D(V_0(a))$ which is not nilpotent on $V_0(a)$, or the conclusion, that no such element exists, if $N_D(V_0(a))$ acts nilpotently on $V_0(a)$. The element a is assumed to be given as a linear combination $a = \sum \alpha_i a_i$ of the basis elements. We suppose that *NON-NILP* runs in polynomial time. We are unable to give such efficient algorithm in general, but, as witnessed by the results of Section 5.2, we have polynomial time algorithms when $D = \mathrm{ad}(\mathcal{L})$, where $\mathcal{L}$ is a Lie algebra over K (K is either finite or a global field) and $V = \mathcal{L}$. We shall also require that $|K| > n$. We denote by Ω a fixed subset of K of size $n+1$. If the characteristic of K is 0, then we insist that $\Omega = \{1, 2, \ldots, n+1\}$. Similarly, if K is a finite extension of $\mathbb{F}_q(X)$ then we assume that Ω consists of nonzero polynomials from $\mathbb{F}_q[X]$ of degree at most $\lceil \log_q(n+2) \rceil$. The main steps of our algorithm are as follows.
1. $a := 0$.
2. $b :=$ *NON-NILP*(a).

3. **if** every element of $N_D(V_0(a))$ is nilpotent on $V_0(a)$ **then return** a.
(At this point $b \in N_D(V_0(a))$ is not nilpotent on $V_0(a)$.)
4. Select an element c from the set $\{a + \alpha(b-a); \alpha \in \Omega\}$ for which $V_0(c)$ is a proper subset of $V_0(a)$.
5. Put $a := c$. If K is infinite, then replace a by an element $\sum \alpha_i a_i$, such that $\alpha_i \in \Omega$ and $\dim_K V_0(\sum \alpha_i a_i) \leq \dim_K V_0(a)$, as provided by Corollary 5.1.3.
6. Go back to step 2.

Upon termination $N_D(V_0(a))$ acts nilpotently on $V_0(a)$, therefore a is then a locally regular element. We claim, that a c can always be selected at step 4.

By construction, W is invariant under the action of the transformations $a + \beta(b-a)$, $\beta \in K$, and hence so is the space V/W. The determinant on V/W of $a + \beta(b-a)$ is a polynomial of β of degree at most $n - \dim_K W$ and not identically zero (take $\beta = 0$). We infer that there exists a subset $\Omega' \subset \Omega$ of size at least $n + 1 - (n - \dim_K W) = \dim_K W + 1$, such that the elements $a + \alpha(b-a)$, $\alpha \in \Omega'$ are nonsingular on the factor V/W. Moreover, as we saw in Section 5.1, the fact that $a + \beta(b-a)$ is nilpotent on W can be described as $h_i(\beta) = 0$, where $h_i(x) \in K[x]$, $1 \leq i \leq \dim_K W$ and $\deg h_i \leq \dim_K W$. The polynomials h_i cannot be all identically zero, because $b = a + 1(b-a)$ is not nilpotent on W. We infer that there is an $\alpha \in \Omega'$ which is not a common zero for all of the polynomials h_i. Then the element $c = a + \alpha(b-a)$ is nonsingular on the factor V/W and non-nilpotent on W, consequently $V_0(c)$ is a proper subset of W. The claim is proved.

The claim ensures, that the quantity $\dim_K V_0(a)$ decreases at each round of the iteration, therefore the body of the loop is executed at most n times. Steps 2–3 require a polynomial number of arithmetical operations in K by assumption on the procedure *NON-NILP*. Step 4 is basic linear algebra, it involves the computation of the images of at most $n+1$ linear transformations. A polynomial bound on step 5 is given in Corollary 5.1.3. Step 5 also ensures, that the coefficients of a and hence b are kept under control. We summarize the result as

Theorem 5.3.1 *Suppose that we are given a subspace $D \leq \mathrm{End}_K V$ ($\dim_K V = n$) by a basis $a_1, a_2, \ldots a_l$. Subject to the assumptions made on K, NON-NILP, and $\Omega \subseteq K$, we can find a locally regular element $a = \sum \alpha_i a_i \in D$ such that the coefficients $\alpha_i \in \Omega$ in deterministic polynomial time.* □

Remark. Suppose that $|K| \geq 2n$, and $\Omega \subseteq K$, $|\Omega| = 2n$. Then an element $a = \sum \alpha_i a_i \in D$ with coefficients α_i drawn uniformly and independently from Ω will be locally regular with probability at least $1/2$. This follows at once from Proposition 5.1.2 and the Sparse

Zeros Lemma. Thus, if K is sufficiently large, then a very efficient randomized method is available.

5.4 Cartan subalgebras

Here we present two algorithms for computing Cartan subalgebras in Lie algebras. Throughout $\mathcal{L}$ denotes a Lie algebra over the field K. We assume that $\mathcal{L}$ is given as input by structure constants with respect to the basis $a_1, a_2, \ldots, a_n$. We have in particular $\dim_K \mathcal{L} = n$. Our objective is to find a (basis of a) Cartan subalgebra $\mathcal{H} \leq \mathcal{L}$.

The first method is based on our algorithm for locally regular endomorphisms.

Theorem 5.4.1 *Suppose that K is either finite, or a global field, and $|K| > n$. Let Ω be a fixed subset of K of size $n+1$. If $\mathrm{char}K = 0$, then we set $\Omega = \{1, 2, \ldots, n+1\}$. If K is a finite extension of $\mathbb{F}_q(X)$ then we assume that Ω consists of nonzero polynomials from $\mathbb{F}_q[X]$ of degree at most $\lceil \log_q(n+2) \rceil$. Then we can find an element $a = \sum \alpha_i a_i \in \mathcal{L}$ with $\alpha_i \in \Omega$ such that $\mathcal{H} := V_0(\mathrm{ad}_{\mathcal{L}} a) \leq \mathcal{L}$ is a Cartan subalgebra of $\mathcal{L}$ in deterministic polynomial time.*

Proof It is known (cf. [Hum], pp. 77–78) that if $a \in \mathcal{L}$ then $\mathcal{H} := V_0(\mathrm{ad}_{\mathcal{L}} a)$ is a subalgebra of $\mathcal{L}$ for which $N_{\mathcal{L}}(\mathcal{H}) = \mathcal{H}$. If in addition $\mathrm{ad}_{\mathcal{L}} a$ is a locally regular endomorphism of $\mathcal{L}$ (as just a linear space over K), then $\mathrm{ad}_{\mathcal{L}} c$ is nilpotent on $\mathcal{L}$ for every element $c \in \mathcal{H}$. By Engel's theorem we obtain that $\mathcal{H}$ is nilpotent, hence a Cartan subalgebra of $\mathcal{L}$. It suffices therefore to find a locally regular element in $D = \mathrm{ad}_{\mathcal{L}}(\mathcal{L})$ with respect to the action on $V = \mathcal{L}$. To this end we can use the algorithm of Theorem 5.3.1. This gives a polynomial time algorithm, because by Theorem 5.2.2 we have an efficient implementation of the procedure *NON-NILP*. $\square$

Concerning our second algorithm, we assume that K is an arbitrary finite field and $\mathcal{L}$ is a *restricted Lie algebra*. Recall (cf. [Jac], pp. 185–194, or [Win], pp. 102–103) that a restricted Lie algebra $\mathcal{L}$ (or a *Lie p-algebra*) over the field K of characteristic $p > 0$ is a Lie algebra over K equipped with a map $x \mapsto x^p$ (called the p-map) such that for every $\alpha \in K$ and $x, y \in \mathcal{A}$

(i) $\mathrm{ad}_{\mathcal{L}}(x^{[p]} = (\mathrm{ad}_{\mathcal{L}} x)^p$;

(ii) $(\alpha x)^{[p]} = \alpha^p x^{[p]}$;

(iii) $(x+y)^{[p]} = x^{[p]} + y^{[p]} + \sum_{i=1}^{p-1} s_i(x,y)$;

where $is_i(x,y)$ is the coefficient of t^{i-1} in $(\mathrm{ad}_{\mathcal{L}}(y - tx))^{p-1}x$. Condition (iii), in paricular the definition of $s_i(x,y)$ is rather thechnical, and can be omitted for Lie algebras with trivial center. Important examples are Lie algebras of associative algebras: if $\mathcal{A}$ is an associative algebra then its Lie algebra $\mathcal{A}_{Lie}$ is a resticted Lie algebra with respect to the natural p-map $x^{[p]} = x^p$.

We only use condition (i). Furthermore, we *do not* require the p-map to be part of the input. The mere existence will be used in the proof of the correctness of the method. For a subset $B \subseteq \mathcal{L}$ we denote by $\mathcal{L}_0(B)$ the subalgebra $\{x \in \mathcal{L} | x\mathrm{ad}_{\mathcal{L}}(b)^n = 0 \text{ for every } b \in B\}$ (the intersection of the Fitting null components of the maps $\mathrm{ad}_{\mathcal{L}}(b)$ for $b \in B$). In the algorithm A stands for a subset of $\mathcal{L}$.

Algorithm Restricted Cartan($\mathcal{L}$):
$A = \emptyset \;\; r := 0.$
forever do
$r := r + 1;$
$a_r :=$Non-nilp($\mathcal{L}_0(A)$)
if $\mathcal{L}_0(A)$ is nilpotent **then return** $\mathcal{L}_0(A)$
$A := A \cup \{a_r\}$
done

Here the auxiliary procedure Non-nilp($\mathcal{S}$) on a subalgebra $\mathcal{S} \leq \mathcal{L}$ returns an element $a \in \mathcal{S}$, such that $\mathrm{ad}_{\mathcal{S}}a$ is not nilpotent, or the conclusion that $\mathcal{S}$ is a nilpotent algebra. By Theorem 5.2.2 this can be done in time polynomial in n and $\log |K|$.

It is immediate that the algorithm terminates in at most n rounds, and that the number of arithmetical operations is bounded by a polynomial of n and $\log |K|$; we have therefore a polynomial time algorithm. It is also clear that upon termination $\mathcal{H} = \mathcal{L}_0(A)$ is a nilpotent subalgebra of $\mathcal{L}$. To show that it is in fact a Cartan subalgebra, we establish that $N_{\mathcal{L}}(\mathcal{H}) = \mathcal{H}$.

Let t be an integer such that $p^t \geq n$. For $a \in \mathcal{L}$ we put $a^* := a^{[p^t]} = (((a^{[p]})^{[p]}) \cdots)^{[p]}$. For a subset $B \subseteq \mathcal{L}$ let $B^* := \{b^* | b \in B\}$. We have $\mathcal{L}_0(B) = \mathcal{L}_0(B^*) = C_{\mathcal{L}}(B^*)$, the centralizer of B^*. The last equality follows because t is sufficiently large to ascertain that $\mathrm{ad}_{\mathcal{L}}(B^*)$ consists of semisimple endomorphisms.

We consider the final state of A. If $A = \emptyset$, then $\mathcal{L}$ is nilpotent, hence a Cartan subalgebra. Otherwise $A = \{a_1, a_2, \ldots, a_r\}$, where $r > 0$. We note that for $1 \leq j \leq r$, a_j is selected from $\mathcal{L}_0(\{a_1 \ldots, a_{j-1}\}) = C_{\mathcal{L}}(\{a_1^*, \ldots, a_{j-1}^*\})$. This implies that if $i \leq j$, then $[a_i^*, a_j] = 0$ and hence $[a_i^*, a_j^*] = [\ldots[[a_i^*, a_j], a_j] \ldots, a_j] = 0$. Taking anti-commutativity into consideration, we have $[a_i^*, a_j^*] = 0$, whenever $1 \leq i, j \leq r$. These facts imply that $A^* \subseteq C_{\mathcal{L}}(A^*) = \mathcal{L}_0(A)$. Now let $y \in N_{\mathcal{L}}(\mathcal{L}_0(A))$. We have $[[y, A^*], A^*] \subseteq [[y, \mathcal{L}_0(A)], A^*] \subseteq [\mathcal{L}_0(A), A^*] = [C_{\mathcal{L}}(A^*), A^*] = 0$. We obtained that $y \in \mathcal{L}_0(A^*) = \mathcal{L}_0(A)$, and hence

$N_{\mathcal{L}}(\mathcal{L}_0(A)) = \mathcal{L}_0(A)$. $\mathcal{L}_0(A)$ is indeed a Cartan subalgebra of $\mathcal{L}$. We have proved the following result:

Theorem 5.4.2 *Let K be a finite field and $\mathcal{L}$ be a restricted Lie algebra over K, given by structure constants. Then the algorithm* Restricted Cartan *finds a Cartan subalgebra $\mathcal{H} \leq \mathcal{L}$ in deterministic polynomial time.* $\square$

5.5 Tori in associative algebras

We conclude this chapter by some remarks on applications to associative algebras. By an n-dimensional *torus* over the field K we understand a commutative and associative separable K-algebra of dimension n. Let $\mathcal{A}$ be a finite dimensional separable associative algebra over the field K. A torus in $\mathcal{A}$ is then a commutative separable subalgebra in $\mathcal{A}$. Let $\mathcal{A}_{Lie}$ denote the Lie algebra of $\mathcal{A}$. Note that if K is of positive characteristic p then $\mathcal{A}_{Lie}$ is a restricted Lie algebra in the natural way. It is easy to see that if $\mathcal{B}$ is a torus in $\mathcal{A}$, then the Fitting zero component of the adjoint action of $\mathcal{B}$ on $\mathcal{A}_{Lie}$ is the centralizer $\mathrm{C}_{\mathcal{A}}(\mathcal{B})$ (cf. [Win], Corollary 4.5.6). Furthermore, it is known (cf. [Win], Theorems 3.2.7.1 and 4.5.17) that Cartan subalgebras of $\mathcal{A}_{Lie}$ are exactly the *maximal tori* in $\mathcal{A}$ (i.e., tori maximal w.r.t. inclusion). The results of the previous section imply

Corollary 5.5.1 *Assume that K is either a finite field or a global field. Given $\mathcal{A}$, a finite dimensional separable K-algebra, we can find a maximal torus in $\mathcal{A}$ in deterministic polynomial time.* $\square$

An element a of the separable K-algebra $\mathcal{A}$ is called a *splitting element* if the K-subalgebra $\langle 1_{\mathcal{A}}, a\rangle$ generated by a and $1_{\mathcal{A}}$ is a maximal torus in $\mathcal{A}$. It is also of interest to find splitting elements in fixed tori. Let $\mathcal{B}$ be an n-dimensional torus over K. Assume that K is large enough, e.g., $K > n^2$. The algorithm will be a straightforward generalization of the textbook technique for finding primitive elements in finite extensions of fields K with $\mathrm{char} K = 0$. In order to demonstrate that the method is in fact a special case of the procedure of Section 5.3, we consider the regular representation of $\mathcal{B}$, although the algorithm will use only computations in $\mathcal{B}$. The regular representation allows us to identify $\mathcal{B}$ with a maximal torus in $\mathcal{A} = \mathrm{M}_n(K)$. We intend to use the general method of Section 5.3 to find a locally regular element in $\mathcal{B}$ with respect to the adjoint action on $V = \mathcal{A}$, ($\mathrm{ad}_{\mathcal{A}} av := av - va$). By [Win], Corollary 4.5.6, we have that the Fitting zero component $V_0(\mathrm{ad}_{\mathcal{A}} a)$ of an element $a \in \mathcal{B}$ is the centralizer of a in $\mathcal{A}$. If follows that for $a, c \in \mathcal{B}$ we have $V_0(\mathrm{ad}_{\mathcal{A}} a) > V_0(\mathrm{ad}_{\mathcal{A}} c)$ if and only if $\langle 1, a\rangle < \langle 1, c\rangle$. Furthermore, the commutativity of $\mathcal{B}$ implies that $V_0(\mathrm{ad}_{\mathcal{A}}a)$ is $\mathcal{B}$-invariant, i.e., $N_{\mathcal{B}}(V_0(\mathrm{ad}_{\mathcal{A}}a)) = \mathcal{B}$ for every $a \in \mathcal{B}$. These facts imply that an element $a \in \mathcal{B}$ is locally regular iff a is a splitting element and suggest the following implementation of the critical steps 2 and 4 in the general procedure:

Step 2. Find an element $b \in \mathcal{B} \setminus \langle 1, a\rangle$ (provided that $\mathcal{B} \neq \langle 1, a\rangle$).

Step 4. Find an element $c \in \{a + \alpha(b - a) | \alpha \in \Omega\}$ such that $\langle 1, c\rangle > \langle 1, b\rangle$.

Both steps can be accomplished in a straightforward way (computing bases of $\langle 1, a\rangle$ and $\langle 1, c\rangle$, respectively). The argument in Section 5.3 implies that this variant of the procedure finds a splitting element in polynomial time, provided that K is either a global field or K is finite with $|K| > n^2$. Note that a more accurate analysis shows that the method already works for $|K| > \binom{n}{2}$. Also note that if the ground field K is small then the Wedderburn decompostion of $\mathcal{B}$ can be found in deterministic polynomial time by the method of [FR] and a straightforward algorithm is available for finding a splitting element in $\mathcal{B}$, provided that such an element exists. We leave the details to the reader.

Corollary 5.5.2 *Assume that K is either a finite field or a global field. Given $\mathcal{A}$, a separable algebra over K, in deterministic polynomial time we can find a splitting element in $\mathcal{A}$, provided that such an element exists.* □

In the literature, several randomized algorithms are based on the technique of splitting elements. It turns out, that in most cases the only step that makes use of randomization is finding a splitting element. Randomized algorithms for finding splitting elements in tori and central simple algebras, based on the Sparse Zeros Lemma, were found by Eberly. The first deterministic polynomial time method for finding splitting elements in central simple algebras over number fields was discovered by Rónyai [Ró5]. We conclude with a list of some important algorithmic problems that can be solved in deterministic polynomial time under the assumption that we have a splitting element at hand.

- Assume that K is either a finite field or a global field. The problem of finding the simple components of a finite dimensional separable K-algebra can be reduced to the task of factoring a single polynomial over K (cf. [Eb1, Eb2]).

- Let K be a number field and $\mathcal{A}$ be a central simple algebra of dimension n^2 over K. A subfield L of $\mathcal{A}$ can be found together with an isomorphism $\phi : L \otimes_K \mathcal{A} \cong \mathrm{M}_n(L)$ (cf. [BR]).

- Let K be a number field and $\mathcal{A}$ be a central simple algebra of dimension n^2 over K. Embeddings $\phi : K \to \mathbb{R}$ can be constucted such that $\mathbb{R} \otimes_{\phi(K)} \mathcal{A} \cong \mathrm{M}_n(\mathbb{R})$. (cf. [Eb3]). We shall use this result in Chapter 6.

Chapter 6

Maximal orders

The results presented in this chapter are based on (partly improved versions of) the methods from the paper [IR], joint work with Lajos Rónyai. In [Ró3], it was shown that the problem of deciding whether the index of a central simple algebra $\mathcal{A}$ over a number field equals a given integer is in $NP \cap coNP$. Actually, it was proved that there exists a maximal $\mathbb{Z}$-order in $\mathcal{A}$ which admits a short description and verification, and the theory of maximal orders and Hasse's principle can be used to determine the index from invariants of maximal orders. To be more specific, the main technical contribution in [Ró3] is a deterministic polynomial time *ff-algorithm* (an algorithm that can make oracle calls to factor integers and polynomials over finite fields) for *testing* maximality of orders. The central result of this chapter is a deterministic polynomial time ff-algorithm for *constructing* maximal orders in semisimple algebras over number fields. A consequence is the existence of a polynomial time ff-algorithm for *computing* the index of central simple algebras. In the important special case of computing the indices of algebras related to representations of finite groups (known as *Schur indices*), the oracle can be substituted with a polynomial time algorithm. The main ideas can be generalized to separable algebras over *global fields*. For algebras over global function fields we obtain *f-algorithms* (that make calls to an oracle for factoring polynomials over the prime field).

We start with the basic definitions. Let R be a Dedekind ring, i.e., a Noetherian integrally closed domain such that every nonzero prime ideal in R is a maximal ideal, K be the field of quotients of R and let $\mathcal{A}$ be a finite dimensional semisimple algebra over K. An *R-order* in $\mathcal{A}$ is a subring Λ of $\mathcal{A}$ satisfying the following properties:

- Λ is a finitely generated module over R,

- Λ has an identity element (this is necessarily the same as the identity element of $\mathcal{A}$ and R),

- Λ generates $\mathcal{A}$ as a linear space over K.

An R-order Λ in $\mathcal{A}$ is a *maximal R-order* if it is not a proper subring of any other R-order of $\mathcal{A}$. It is known that if $\mathcal{A}$ is a separable commutative K-algebra (e.g., $\mathcal{A}$ is a finite separable extension field of K), then the integral closure Λ of R in $\mathcal{A}$ defined by

$$\Lambda = \{x \in \mathcal{A} | \text{there exists a monic polynomial } f(X) \in R[X] \text{ such that } f(x) = 0\}$$

is the a unique maximal order in $\mathcal{A}$.

If P is a maximal ideal in R then we can consider the localization R_P of R at P:

$$R_P = \{\frac{\alpha}{\beta} | \alpha \in R, \beta \in R \setminus P\}.$$

R_P is a discrete valuation ring, i.e., a local Dedekind ring. We shall represent R_P-orders with R-orders. If Λ is an R-order then the localization

$$\Lambda_P = R_P \Lambda = \{\frac{1}{\beta} x | x \in \Lambda,\ \beta \in R \setminus P\}$$

is an R_P-order. We say that Λ is *locally maximal* at P if Λ_P is a maximal R_P-order. It turns out that Λ is maximal if and only if it is locally maximal at every maximal ideal P of R.

Let $\mathcal{A}$ be a semisimple algebra over the number field K, and D denote the ring of integers in K. The key result of this chapter is a polynomial time f-algorithm that, given $\mathcal{A}$ and a rational prime p, constructs a $\mathbb{Z}$-order Λ which is locally maximal at the prime p. An application is a polynomial time ff-algorithm for finding a (globally) maximal D-order in $\mathcal{A}$, as well as one for computing the index of $\mathcal{A}$. For the analogous problems over global function fields, these ideas lead to polynomial time f-algorithms (making oracle calls to factor polynomials over the prime field).

As an application to representation theory of finite groups, we give a deterministic polynomial time method to compute the degrees of the irreducible submodules of a finite matrix group over an algebraic number field K of a finite group G given by generators.

The organization of this chapter is as follows. Section 6.1 contains the basic statements from the theory of orders we need. Most of the material can be found in [Re]. Proofs are given only where a precise reference was hard to locate.

In Section 6.2 we collect some statements about the radicals of orders over discrete valuation rings. These play an important role in the study of *extremal orders* later on.

Section 6.3 and in particular Proposition 6.3.1 and Theorem 6.3.5 contain the statements which serve as the theoretical foundation for our algorithms. These two statements enable us to reduce the problem of finding maximal orders over discrete valuation rings to that of decomposing associative algebras over the residue class fields. The ideas presented here are not new. They were used by Jacobinski (see [Ja] or [Re], Chapter 39) in his approach to the theory of hereditary orders. We include proofs because Jacobinski worked with complete local rings. In the statements here the completeness of R is not assumed. Also, largely due to the fact that weaker results are sufficient for our purposes, it was possible to simplify some of the original arguments.

Section 6.4 contains the algorithms for computing maximal orders. Theorem 6.4.1 provides the basic 'iteration step' of our subsequent methods for constructing (locally) maximal orders. We describe an algorithm that for a given order Λ constructs an order Γ containing Λ such that Γ is "locally greater" than Λ if such an order exists.

In Corollary 6.4.3 we give a polynomial time ff-algorithm for constructing a maximal D-order Λ in a semisimple algebra $\mathcal{A}$ over an algebraic number field K. This settles in the affirmative the question raised in [Ró3].

In Section 6.5 algorithms for computing the index are presented. Perhaps the most interesting result of this chapter is Theorem 6.5.5. We propose a deterministic polynomial time algorithm to compute the dimensions of the irreducible G-submodules of K^n, where K is a number field and G is a finite subgroup of $\mathrm{GL}_n(K)$ given by generators. An important special case is computing the degrees of the irreducible representations over an algebraic number field K of a finite group G given by a multiplication table.

Notation and terminology

Throughout this chapter we keep ourselves to the following notation and terminology:

- R: a Dedekind ring, i.e., a Noetherian integrally closed domain in which the nonzero prime ideals are maximal

- K: the field of quotients of R

- P: a unique maximal ideal of R

- π: a prime element of R in case of R is a principal ideal domain, i.e., an element such that the ideal $P = (\pi)$ is maximal. Typically, π stands for either a rational prime p (case $R = \mathbb{Z}$) or an irreducible polynomial $f(X) \in \mathbb{F}_q[X]$ (case $R = \mathbb{F}_q[X]$).

- $\mathcal{A}$: a finite dimensional separable algebra over K

- *order*: we use this term for an R-order in the K-algebra $\mathcal{A}$.

- Γ, Λ: orders

- *lattice*: If V is a finite dimensional K-space, an R-lattice in V is a finitely generated submodule M of the R-module V. If, in addition, M is a K-space generating set of the entire space V, we say that M is a *full R-lattice* in V. Full lattices in the vector space $\mathcal{A}$ are of particular interest. Orders are special cases of full lattices in $\mathcal{A}$.

- *radical*: the Jacobson radical of a ring or algebra, denoted by $\mathrm{Rad}(R)$, $\mathrm{Rad}(\Lambda)$, etc.

6.1 Basic facts about orders

In this section we collect the basic facts and some elementary results from the theory of orders we need later on. In this section we assume that $\mathcal{A}$ is a separable algebra, i.e., semisimple and the centers of its simple components are separable extensions of the ground field K.

Reduced trace forms and discriminants

First we introduce the *reduced trace* function of a semisimple algebra using a sequence of progressively more general definitions (for a central simple algebra, then a simple algebra, and finally for a semisimple algebra).

We start from the trace of the left regular representation of $\mathcal{A}$. To be more specific, the *trace* $\mathrm{Tr}_{\mathcal{A}/K}(x)$ of an element $x \in \mathcal{A}$ over K is the trace of the the K-linear transformation $L_x : \mathcal{A} \to \mathcal{A}$ defined by $L_x(a) = xa$ for $a \in \mathcal{A}$.

If $\mathcal{A}$ is a full matrix algebra over the field E, $\dim_E \mathcal{A} = n^2$, then there is another way to define traces of elements of $\mathcal{A}$. Namely, if we have an isomorphism $\phi\colon \mathcal{A} \cong \mathrm{M}_n(E)$ then we can take $\mathrm{tr}_{\mathcal{A}/E}(x)$ as the trace of the matrix ϕx. This is independent from the choice of the isomorphism ϕ.

If $\mathcal{A}$ is a central simple K-algebra, $\dim_K \mathcal{A} = n^2$, then there exists an extension field E of K which splits $\mathcal{A}$, i.e., $E \otimes_K \mathcal{A} \cong \mathrm{M}_n(E)$. It can be shown that $\mathrm{tr}_{\mathcal{A}/K}(x) := \mathrm{tr}_{E\otimes_K \mathcal{A}/E}(x \otimes 1) \in K$ is independent of the choice of the splitting field E and we have $n\mathrm{tr}_{\mathcal{A}/K}(x) = \mathrm{Tr}_{\mathcal{A}/K}(x)$. Consequently if the characteristic of K is zero (or prime to n), then $tr_{\mathcal{A}/K}(x) = \frac{1}{n}\mathrm{Tr}_{\mathcal{A}/K}(x)$.

If $\mathcal{A}$ is a simple K-algebra with center L then we can take $\mathrm{tr}_{\mathcal{A}/K}(x) := \mathrm{Tr}_{L/K}\mathrm{tr}_{\mathcal{A}/L}(x)$. If $\mathcal{A}$ is a semisimple K-algebra with Wedderburn decomposition $\mathcal{A} = \mathcal{A}_1 \oplus \ldots \oplus \mathcal{A}_r$, then we can define $\mathrm{tr}_{\mathcal{A}/K}(x) := \mathrm{tr}_{\mathcal{A}_1/K}(x_1) + \ldots + \mathrm{tr}_{\mathcal{A}_r/K}(x_r)$, where x_i is the image of x under the projection $\mathcal{A} \to \mathcal{A}_i$ onto the ith simple component of $\mathcal{A}$. We call $\mathrm{tr}_{\mathcal{A}/K}(x)$ the *reduced trace* of x over K. The map $\langle\ ,\ \rangle\colon \mathcal{A}\times\mathcal{A} \to K$ defined by $\langle x, y\rangle := \mathrm{tr}_{\mathcal{A}/K}(xy)$ is a K-bilinear function and is called the bilinear trace form of $\mathcal{A}$ over K. If $\mathcal{A}$ is *separable* over K then

$\langle\ ,\ \rangle$ is a nondegenerate bilinear form. For the rest of this subsection we assume that $\mathcal{A}$ is separable over K.

We shall omit the subscript $\mathcal{A}/K$ from $\mathrm{Tr}_{\mathcal{A}/K}$ and $\mathrm{tr}_{\mathcal{A}/K}$ whenever $\mathcal{A}$ and K are clear from the context.

Let Λ be an R-order in $\mathcal{A}$. Then for every element $x \in \Lambda$, we have $\mathrm{tr}(x) \in R$ (cf. [Re], Thm. 10.1). Let $n = \dim_K \mathcal{A}$. The *discriminant* of the order Λ is the ideal $\mathrm{disc}(\Lambda)$ in R generated by the set

$$\{\det\left(\mathrm{tr}(x_i x_j)\right)_{i,j=1}^n \mid (x_1, \ldots, x_n) \in \Lambda^n\}.$$

Proposition 6.1.1 *Assume that* $\Lambda \subseteq \Gamma$. *Then* $\mathrm{disc}(\Gamma) \supseteq \mathrm{disc}(\Lambda)$ *and* $\Lambda = \Gamma$ *if and only if* $\mathrm{disc}(\Gamma) = \mathrm{disc}(\Lambda)$.

Proof [Re], Exercise 10.3 or 4.13. □

From a generating set of Λ as an R-module we can easily obtain a nonzero multiple of $\mathrm{disc}(\Lambda)$: we select a subset $\{x_1, \ldots, x_n\}$ of the generating set which is a K-basis of $\mathcal{A}$.

Proposition 6.1.2 *Let* $\{x_1, \ldots, x_n\} \subseteq \Lambda$ *be a* K*-basis of* $\mathcal{A}$. *Then the principal ideal generated by the nonzero determinant* $d = \det\left(\mathrm{tr}(x_i x_j)\right)_{i,j=1}^n$ *is contained in the discriminant.*

Proof Obvious. □

Proposition 6.1.3 *Let* $\{x_1, \ldots, x_n\}$ *and* d *be as in Proposition 6.1.2. Let* Γ *be any order containing* Λ. *Then* $d\Gamma \subseteq R\{x_1, \ldots, x_n\} \subseteq \Lambda$.

Proof The proof is a version of the argument given in [Re], Thm. 10.3. Let $a \in \Gamma$ and put $a = \sum \gamma_i a_i$, where $\gamma_i \in K$. Then

$$\mathrm{tr}(aa_j) = \sum_{i=1}^{n} \gamma_i \mathrm{tr}(a_i a_j), \quad 1 \leq j \leq n.$$

We have $\mathrm{tr}(aa_j), \mathrm{tr}(a_i a_j) \in R$ because $aa_j, a_i a_j \in \Gamma$ and therefore they are integral elements over R. If we use Cramer's rule to solve the system of linear equations above for the γ_i, we obtain that $\gamma_i = \frac{\alpha_i}{d}$ for some $\alpha_i \in R$. □

Note that if R is a principal ideal domain then every R-order Λ admits an R-basis, say $\{x_1, \ldots, x_n\}$, and the discriminant $\mathrm{disc}(\Lambda)$ is the principal ideal generated by the determinant $\det\left(\mathrm{tr}(x_i x_j)\right)_{i,j=1}^n$ (cf. [Re], Theorem 10.2).

Localizations

If R is a Dedekind domain with quotient field K and P is a prime ideal in R then the ring of quotients $R_P = (R \setminus P)^{-1}R \subset K$ is a discrete valuation ring. For an R-lattice M in $\mathcal{A}$ we can define the localization at P as follows: $M_P = R_P M \subset \mathcal{A}$. M_P is an R_P-lattice. If M is a full R-lattice in $\mathcal{A}$ (i.e., $KM = \mathcal{A}$), then M_P is a full R_P-lattice in $\mathcal{A}$. If Λ is an R-order then Λ_P is an R_P-order, moreover Λ is a maximal R-order if and only if Λ_P is a maximal R_P order for every prime ideal P of R. More generally,

Proposition 6.1.4 *If Γ and Λ are R-orders in $\mathcal{A}$ such that $\Lambda \subset \Gamma$ then there exist a prime ideal P of R such that $\Lambda_P \subset \Gamma_P$.*

Proof [Re], Theorem 3.15. □

The next statement demonstrates a simple but useful connection between the orders Λ and Λ_P.

Proposition 6.1.5 *Let Λ be an R-order in $\mathcal{A}$. The map $\phi : x \mapsto x + P\Lambda_P$ $(x \in \Lambda)$ induces an isomorphism of rings $\Lambda/P\Lambda \cong \Lambda_P/P\Lambda_P$.*

Proof Clearly $\phi : \Lambda \to \Lambda_P/P\Lambda_P$ is an epimorphism of rings. It is straightforward to check that $ker(\phi) = P\Lambda$. □

To be more specific, if R happens to be a principal ideal domain, and π is a prime element in R, i.e., the principal ideal (π) is a maximal ideal in R, then we can write

$$R_{(\pi)} = \{r/s \in K | r, s \in R,\ \gcd(\pi, s) = 1\}.$$

$R_{(\pi)}$ is a discrete valuation ring with unique maximal ideal $R_{(\pi)}\pi$. If Λ is an R-order, then we use the notation $\Lambda_{(\pi)} = R_{(\pi)}\Lambda$.

Important examples are when $R = \mathbb{Z}$ and $\pi = p$, a rational prime, or $R = \mathbb{F}_q[X]$ and $\pi = f(X)$, a (monic) irreducible polynomial in $\mathbb{F}_q[X]$, respectively. Then we use the notation $\mathbb{Z}_{(p)}$, $\Lambda_{(p)}$, $\mathbb{F}_q[X]_{(f(X))}$, and $\Lambda_{(f(X))}$, respectively.

Orders over extensions

Assume that L is a finite separable extension of K, and $\mathcal{A}$ is finite dimensional separable L-algebra. Let D be the integral closure of R in L. The next statement will be useful when we change the ring of coefficients from R to D.

Lemma 6.1.6 *Let L be a finite separable extension field of K, $\mathcal{A}$ a finite dimensional separable algebra over L and let Λ be an R-order in $\mathcal{A}$. Let D be the integral closure of R in L. Then $\Gamma = D\Lambda$ is a D-order containing Λ. As a consequence, if Λ is a maximal R-order then $\Lambda \cap L (= \Lambda \cap L1_{\mathcal{A}}) = D$. Moreover, a maximal R-order in $\mathcal{A}$ is a maximal D-order as well.*

Proof It is straightforward to check that $D\Lambda$ (the finite sums of the form $\sum \alpha_i x_i$, $\alpha_i \in D$, $x_i \in \Lambda$) is a ring which is a finitely generated D-module. Also we have $1_\Lambda \in D\Lambda$, therefore Γ is indeed a D-order. Now the rest is obvious. $\square$

We will use the following statement to compute local properties of D-orders without explicitly computing D.

Lemma 6.1.7 *Let L be a finite separable extension field of K, $\mathcal{A}$ a finite dimensional separable algebra over L and let D be the integral closure of R in L. Assume that Λ is an R-order in $\mathcal{A}$ such that Λ is locally maximal at the prime ideal P of R. Then the R-order $\Gamma = D\Lambda$ is a D-order as well, and Γ is locally maximal at every prime ideal $\mathcal{P}$ of D above P. Let $\mathcal{B} = \Lambda/P\Lambda$ and $\mathcal{C}$ be the image of $\Lambda \cap L$ at the natural map $\phi : \Lambda \to \mathcal{B}$. The map $\mathcal{P} \mapsto \phi(\mathcal{P} \cap \Lambda)$ is a bijection between the prime ideals of D and the maximal ideals of $\mathcal{C}$ such that $\Gamma/\mathcal{P}\Gamma \cong \mathcal{B}/\phi(\mathcal{P} \cap \Lambda)\mathcal{B}$.*

Proof The local maximality of Λ implies $\Gamma_P = \Lambda_P$. Let $\Delta = \Lambda \cap L$. Then $\Delta_P = \Lambda_P \cap L = \Gamma_P \cap L = (\Gamma \cap L)_P = D_P$, i.e., Δ is also locally maximal at P (in L). Let $\psi : x \mapsto x + P\Gamma_P$ be the natural map $\Gamma_P \to \Gamma/P\Gamma_P$. By Proposition 6.1.5, the restriction of ψ to Γ induces an isomorphism $\Gamma/P\Gamma \cong \Gamma_P/P\Gamma_P$. Similarly, the restriction of ψ to Λ induces an isomorphism $\Lambda/P\Lambda \cong \Gamma_P/P\Gamma_P$. In fact, we can identify $\mathcal{B}$ with $\Gamma_P/P\Gamma_P$ and ϕ with the restriction of ψ to Λ. Using these identifications, we have $\psi(D) = \psi(D_P) = \psi(\Delta_P) = \psi(\Delta) = \mathcal{C}$. The kernel of the restriction of ψ to D_P is PD_P. It follows that ψ induces a bijection between the ideals of D containing PD and the ideals of $\mathcal{C}$. For every maximal ideal $\mathcal{P}$ of D above P, we have $\psi(\mathcal{P}\Gamma)) = \psi(\mathcal{P})\psi\Gamma = \psi(\mathcal{P})\mathcal{B}$, whence ψ induces an isomorphism $\Gamma/\mathcal{P}\Gamma \cong \mathcal{B}/\psi(\mathcal{P} \cap \Lambda)\mathcal{B}$. $\square$

Orders over $\mathbb{Z}$ and $\mathbb{F}_q[X]$

There are some simple examples of orders. If M is a full R-lattice in $\mathcal{A}$ (i.e., $KM = \mathcal{A}$) then the left order of M defined by $\mathcal{O}_l(M) = \{x \in \mathcal{A} \mid xM \subseteq M\}$ is an R-order in $\mathcal{A}$ (cf. p. 109, [Re]). The right order is defined in a similar way. These examples offer an important algorithmic tool for constructing orders.

For the rest of this section we assume that $R = \mathbb{Z}$ or $R = \mathbb{F}_q[X]$. The following statement is an easy generalization of [Ró3], Theorem 3.1.

Proposition 6.1.8 *If a full R-lattice M in the K-algebra $\mathcal{A}$ is given by an R-basis then $\mathcal{O}_l(M)$ has an R-basis of size* $(\mathrm{size}(\mathcal{A}) + \mathrm{size}(M))^{O(1)}$ *and such a basis can be computed in time* $(\mathrm{size}(\mathcal{A}) + \mathrm{size}(M))^{O(1)}$.

Proof Let $b_1, \dots, b_n$ be the given R-basis of M. Since $KM = \mathcal{A}$, we can express $1_\mathcal{A}$ as a K-linear combination of $b_1, \dots, b_n$ via solving a system of linear equations over K. Computing a common denominator leads to finding $s \in R$ in time $(\mathrm{size}(\mathcal{A}) + \mathrm{size}(M))^{O(1)}$ such that $s1_\mathcal{A} \in M$. For such an s we have

$$\mathcal{O}_l(M) = \{x \in s^{-1}M | xM \subseteq M\}.$$

Finding an R-basis of $\mathcal{O}_l(M)$ in terms of $s^{-1}b_1, \dots, s^{-1}b_n$ is equivalent to computing an R-basis of the R-integral solutions of a system of linear equations. This can also be done in polynomial time. (For the case $R = \mathbb{Z}$, the reader is referred to [Fr] or [KB]. Algorithms for $R = \mathbb{F}_q[X]$ are outlined in Section 1.2.) $\square$

The following statement gives a tool to reduce the problem of enlarging an R-order to a similar problem for $R_{(\pi)}$-orders.

Lemma 6.1.9 *Let π be a prime element in R and Γ be an R-order. Suppose that $\mathcal{J}$ is an ideal of $\Gamma_{(\pi)}$ such that $\mathcal{J} \supseteq \pi\Gamma_{(\pi)}$ and $\mathcal{O}_l(\mathcal{J}) \supset \Gamma_{(\pi)}$. Let $\mathcal{I}$ denote the inverse image of $\mathcal{J}$ with respect to the embedding $\Gamma \to \Gamma_{(\pi)}$. Then we have $\mathcal{I} \supseteq \pi\Gamma$, $\mathcal{O}_l(\mathcal{I}) \supset \Gamma$, and $\mathcal{O}_l(\mathcal{I})_{(\pi)} \supset \Gamma_{(\pi)}$.*

Proof Clearly $\mathcal{I} \supseteq \pi\Gamma$ and $\mathcal{I}$ is an ideal of Γ. Let $a_1, a_2, \dots, a_t$ be a generating set of $\mathcal{I}$, as an R-module. Then the images of the elements a_i (which will also be denoted by a_i) form a generating set of $\mathcal{J}$ as an $R_{(\pi)}$-module. Now let $a \in \mathcal{O}_l(\mathcal{J}) \setminus \Gamma_{(\pi)}$. Then for $i = 1, \dots, t$ we have

$$aa_i = \frac{\alpha_{i1}}{\beta_{i1}}a_1 + \frac{\alpha_{i2}}{\beta_{i2}}a_2 + \dots + \frac{\alpha_{it}}{\beta_{it}}a_t,$$

where $\alpha_{ij}, \beta_{ij} \in R$ and π does not divide β_{ij}. Now put $\beta = \prod_{i,j} \beta_{ij}$. Then it is straightforward to check that $\beta a\mathcal{I} \subseteq \mathcal{I}$ and consequently $\beta a \in \mathcal{O}_l(\mathcal{I})$. Finally we observe that $\beta a \notin \Gamma_{(\pi)}$, for otherwise we have $a \in \Gamma_{(\pi)}$. The proof is complete. $\square$

The next statement provides a bound on the number of iterations in algorithms which successively increase orders until a maximal order is obtained.

Proposition 6.1.10 *Assume that we have the strictly increasing chain $\Lambda_0 \subset \ldots \subset \Lambda_m$ of R-orders in $\mathcal{A}$. Let $d_i \in R$ be a generator of the ideal $\mathrm{disc}(\Lambda_i)$, for $i = 0, \ldots, m$. In the case $d \in R = \mathbb{Z}$, let $|d|$ denote the usual absolute value of the integer d, while in the case $0 \neq d = d(X) \in R = \mathbb{F}_q[X]$, let $|d| = 2^{\deg d(X)}$. Then*

$$m \leq \frac{1}{2} \log_2 |d_0/d_m| \leq \frac{1}{2} \log_2 |d_0|.$$

Proof For each $i = 0, \ldots, m-1$, $|\frac{d_i}{d_{i+1}}| > 1$ is the square of an integer ($|\det T_i|^2$, where T_i is a matrix transforming an R-basis of Λ_{i+1} to an R-basis of Λ_i). We obtain the statement by taking logarithm of

$$|\frac{d_0}{d_m}| = \prod_{i=0}^{m-1} |\frac{d_i}{d_{i+1}}| \geq 2^{2m}.$$

□

6.2 Radicals of orders over local rings

First we recall some basic facts about the Jacobson radical of rings. For proofs, see for example [Re], Section 6a. Let $\mathcal{S}$ denote an arbitrary ring with an identity element. $\mathrm{Rad}(\mathcal{S})$, the Jacobson radical of $\mathcal{S}$, is the set of elements $x \in \mathcal{S}$ such that $xM = (0)$ for all simple left (or, equivalently, $Mx = (0)$ for all simple right) modules M over $\mathcal{S}$. $\mathrm{Rad}(\mathcal{S})$ is a two-sided ideal in $\mathcal{S}$ containing every nilpotent one-sided ideal of $\mathcal{S}$. Also, $\mathrm{Rad}(\mathcal{S})$ can be characterized as the intersection of the maximal left ideals in $\mathcal{S}$, and, equivalently, as the intersection of the maximal right ideals in $\mathcal{S}$. If $\mathcal{S}$ is left or right Artinian (this holds for example if $\mathcal{S}$ is a finite dimensional algebra with identity over a field) then $\mathrm{Rad}(\mathcal{S})$ is the maximal nilpotent ideal in $\mathcal{S}$.

After these preliminaries let us return to our rings of interest. We assume that R is a discrete valuation ring, P is the unique nonzero prime ideal of R, K is the field of quotients of R, and Λ is an R-order in a finite dimensional semisimple K-algebra $\mathcal{A}$.

Proposition 6.2.1 *The residue class ring $\bar{\Lambda} = \Lambda/P\Lambda$ is an algebra with identity element over the residue class field $\bar{R} = R/P$ and $\dim_K \mathcal{A} = \dim_{\bar{R}} \bar{\Lambda}$. If $\phi\colon \Lambda \to \bar{\Lambda}$ is the canonical epimorphism, then $P\Lambda \subseteq \mathrm{Rad}(\Lambda) = \phi^{-1}\mathrm{Rad}(\bar{\Lambda})$ and ϕ induces a ring isomorphism $\Lambda/\mathrm{Rad}(\Lambda) \cong \bar{\Lambda}/\mathrm{Rad}(\bar{\Lambda})$. As a consequence, a left (or right) ideal $\mathcal{I}$ of Λ is contained in $\mathrm{Rad}(\Lambda)$ if and only if $\mathcal{I}$ is nilpotent modulo $P\Lambda$, i.e., there exists a positive integer t such that $\mathcal{I}^t \subseteq P\Lambda$.*

Proof Most of the statements are proved in [Re], Thm. 6.15. The claim about the dimensions follows directly from the fact that R is a principal ideal ring and Λ is a free R-module. As for the 'only if' part of the last statement, every nilpotent ideal of $\bar{\Lambda}$ is contained in $\mathrm{Rad}(\bar{\Lambda})$. □

Proposition 6.2.2 *If $\Lambda \subseteq \Gamma$ are R-orders, then there exists a positive integer s such that* $\mathrm{Rad}(\Gamma)^s \subseteq \Lambda$. *For any such s,* $\mathrm{Rad}(\Gamma)^s \subseteq \mathrm{Rad}(\Lambda)$ *is an ideal in* Λ.

Proof (A part of the argument below can be found in [Re], hint to Ex. 39.3.)
Using that $\Gamma \supseteq \Lambda$ are full R-modules in $\mathcal{A}$ over a discrete valuation ring R, from Proposition 6.2.1 we infer that there exist positive integers u and t such that $P^u\Gamma \subseteq \Lambda$ and $\mathrm{Rad}(\Gamma)^t \subseteq P\Gamma$. Now $s = tu$ will suffice to prove the first claim. If for some s we have $\mathcal{I} = \mathrm{Rad}(\Gamma)^s \subseteq \Lambda$ then $\mathcal{I}$ is an ideal in Λ because $\Lambda\mathcal{I} \subseteq \Gamma\mathcal{I} = \mathcal{I}$ and $\mathcal{I}\Lambda \subseteq \mathcal{I}\Gamma = \mathcal{I}$. Finally for the integers t and u we have

$$\mathcal{I}^{t(u+1)} = \mathrm{Rad}(\Gamma)^{st(u+1)} \subseteq (P\Gamma)^{s(u+1)} \subseteq (P\Gamma)^{(u+1)} = P^{(u+1)}\Gamma = P \cdot P^u\Gamma \subseteq P\Lambda.$$

Proposition 6.2.1 implies that $\mathcal{I} \subseteq \mathrm{Rad}(\Lambda)$. □

The following observation plays an important role in Jacobinski's approach (cf. [Ja]) to hereditary orders.

Proposition 6.2.3 *Let $\Lambda \subseteq \Gamma$ be R-orders in $\mathcal{A}$ such that* $\mathrm{Rad}(\Gamma) \subseteq \Lambda$. *Then for any order Λ' such that $\Lambda \subseteq \Lambda' \subseteq \Gamma$ we have* $\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda')$. *The canonical map $\phi\colon \Gamma \to \bar{\Gamma} =$* $\Gamma/\mathrm{Rad}(\Gamma)$ *induces a bijection $\Lambda' \mapsto \Lambda'/\mathrm{Rad}(\Gamma)$ between the set of orders Λ' lying between Λ and Γ and the set of the subalgebras of the R/P-algebra $\bar{\Gamma}$ containing* $\Lambda/\mathrm{Rad}(\Gamma)$. *Moreover if $\Lambda \subseteq \Lambda' \subseteq \Gamma$, then we have* $\mathrm{Rad}(\Lambda') = \phi^{-1}\mathrm{Rad}(\phi\Lambda')$.

Proof We have $\mathrm{Rad}(\Gamma) \subseteq \Lambda \subseteq \Lambda'$. From this Proposition 6.2.2 implies that $\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda')$. The statement about the correspondence of R-orders and R/P-subalgebras is obvious once we observe that any R-subalgebra Λ' such that $\Lambda \subseteq \Lambda' \subseteq \Gamma$ is actually an R-order. As for the last statement, we note that if $\mathcal{J}$ is a maximal left ideal of Λ' then $\mathrm{Rad}(\Gamma) \subseteq \mathcal{J}$, because $\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda')$. We infer that ϕ induces a bijection between the set of the maximal left ideals of Λ' and the set of the maximal left ideals of $\Lambda'/\mathrm{Rad}(\Gamma)$ and the statement follows. □

6.3 Extremal orders

In this section R is a discrete valuation ring. For R-orders in $\mathcal{A}$ we introduce the following partial ordering: Γ *radically contains* Λ if and only if $\Gamma \supseteq \Lambda$ and $\mathrm{Rad}(\Gamma) \supseteq \mathrm{Rad}(\Lambda)$. The orders maximal with respect to this partial ordering are called *extremal*. Maximal orders are obviously extremal. The notion of extremal orders has been introduced in [Ja]. The next statement is from [Ja], Proposition 1. We note first that if Λ is an R-order then $P\Lambda \subseteq \mathrm{Rad}(\Lambda)$, so that $\mathrm{Rad}(\Lambda)$ is a full R-lattice. Therefore $\mathcal{O}_l(\mathrm{Rad}(\Lambda))$ is an R-order.

Proposition 6.3.1 *For any R-order Λ, the order $\mathcal{O}_l(\mathrm{Rad}(\Lambda))$ radically contains Λ. Moreover, an R-order Λ of $\mathcal{A}$ is extremal if and only if $\Lambda = \mathcal{O}_l(\mathrm{Rad}(\Lambda))$ (if and only if $\Lambda = \mathcal{O}_r(\mathrm{Rad}(\Lambda))$).*

Proof Since $\mathrm{Rad}(\Lambda)$ is an ideal in Λ, $\Lambda \subseteq \mathcal{O}_l(\mathrm{Rad}(\Lambda))$. Also, $\mathrm{Rad}(\Lambda)$ is a left ideal in $\mathcal{O}_l(\mathrm{Rad}(\Lambda))$ and by Proposition 6.2.1 for some t we have

$$\mathrm{Rad}(\Lambda)^t \subseteq P\Lambda \subseteq P\mathcal{O}_l(\mathrm{Rad}(\Lambda)),$$

hence

$$\mathrm{Rad}(\Lambda) \subseteq \mathrm{Rad}(\mathcal{O}_l(\mathrm{Rad}(\Lambda))).$$

This implies that $\mathcal{O}_l(\mathrm{Rad}(\Lambda))$ radically contains Λ. We infer that if Λ is extremal then $\Lambda = \mathcal{O}_l(\mathrm{Rad}(\Lambda))$.

In the other direction, we suppose that $\Lambda = \mathcal{O}_l(\mathrm{Rad}(\Lambda))$ and Γ is an order radically containing Λ. By Proposition 6.2.2 there exists an integer $s > 0$ such that $\mathrm{Rad}(\Gamma)^s \subseteq \mathrm{Rad}(\Lambda)$. For any $s > 1$ with this property we have

$$\mathrm{Rad}(\Gamma)^{s-1}\mathrm{Rad}(\Lambda) \subseteq \mathrm{Rad}(\Gamma)^{s-1}\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda),$$

implying that

$$\mathrm{Rad}(\Gamma)^{s-1} \subseteq \mathcal{O}_l(\mathrm{Rad}(\Lambda)) = \Lambda.$$

Proposition 6.2.2 implies that $\mathrm{Rad}(\Gamma)^{s-1} \subseteq \mathrm{Rad}(\Lambda)$. Continuing in this way we obtain $\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda)$ and consequently $\mathrm{Rad}(\Gamma) = \mathrm{Rad}(\Lambda)$. We conclude that

$$\Gamma \subseteq \mathcal{O}_l(\mathrm{Rad}(\Gamma)) = \mathcal{O}_l(\mathrm{Rad}(\Lambda)) = \Lambda$$

and $\Gamma = \Lambda$. $\square$

Proposition 6.3.2 *Assume that $\Lambda \subseteq \Gamma$ are R-orders. Then $\Lambda + \mathrm{Rad}(\Gamma)$ is an R-order radically containing Λ.*

Proof It is straightforward to verify that $\Lambda' = \Lambda + \mathrm{Rad}(\Gamma)$ is an R-order containing Λ. Next, using the characterization of radical-ideals from Proposition 6.2.1, we obtain that $\mathrm{Rad}(\Lambda) + \mathrm{Rad}(\Gamma)$ is an ideal of Λ' and $\mathrm{Rad}(\Lambda) + \mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda')$. □

Proposition 6.3.3 *Let $\Lambda \subseteq \Gamma$ be R-orders and suppose that Λ is extremal. Then $\mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda)$.*

Proof An immediate consequence of Propositions 6.3.2 and 6.2.2. □

We remark that if Λ is an R-order in $\mathcal{A}$ such that $\mathrm{Rad}(\Lambda) = P\Lambda = \pi\Lambda$ then Λ is a maximal order. Indeed, $\mathcal{O}_l(\pi\Lambda) = \mathcal{O}_l(\Lambda) = \Lambda$, hence Λ is extremal by Proposition 6.3.1. If $\Gamma \supseteq \Lambda$ then by Proposition 6.3.3 we have

$$\pi\Gamma \subseteq \mathrm{Rad}(\Gamma) \subseteq \mathrm{Rad}(\Lambda) = \pi\Lambda$$

implying that $\pi\Gamma = \pi\Lambda$ and $\Gamma = \Lambda$.

Theorem 6.3.5 plays a key role in our method for constructing a maximal R-order. The statement and the proof is a simplified version of [Ja], Proposition 2. We need first an auxiliary lemma on semisimple algebras.

Lemma 6.3.4 *Let $\mathcal{B}$ be a finite dimensional semisimple algebra over a field F. Let $\mathcal{C}$ be a maximal subalgebra of $\mathcal{B}$, such that $\mathrm{Rad}(\mathcal{C}) \neq 0$. Then there exists a two-sided ideal $\mathcal{J}$ of $\mathcal{C}$ minimal among those containing $\mathrm{Rad}(\mathcal{C})$ which is a left ideal of $\mathcal{B}$.*

Proof First we reduce the statement to the special case when $\mathcal{B}$ is simple. In general by Wedderburn's theorem we have $\mathcal{B} = \mathcal{B}_1 \oplus \cdots \oplus \mathcal{B}_r$, where the direct summands $\mathcal{B}_i$ are simple algebras. We observe first, that $\mathcal{C}$ contains the center $C(\mathcal{B})$ of $\mathcal{B}$. Indeed, for the algebra $\mathcal{C}' = < \mathcal{C}, C(\mathcal{B}) >$ we have $\mathcal{C}' \supseteq \mathcal{C}$. Also, it is straightforward to verify that an element $0 \neq c \in \mathrm{Rad}(\mathcal{C})$ generates a nilpotent left ideal in $\mathcal{C}'$ as well, therefore $\mathrm{Rad}(\mathcal{C}') \neq 0$. This implies that $\mathcal{C}' \subset \mathcal{B}$ and hence $\mathcal{C}' = \mathcal{C}$ and $\mathcal{C} \supseteq C(\mathcal{B})$.

We infer that $\mathcal{C}$ contains the identity elements $e_i \in \mathcal{B}_i$ of the ideals $\mathcal{B}_i$ and consequently we have $\mathcal{C} = e_1\mathcal{C} \oplus \cdots \oplus e_r\mathcal{C}$. Now the maximality of $\mathcal{C}$ implies the existence of an index i, such that $e_i\mathcal{C}$ is a maximal subalgebra of the simple algebra $\mathcal{B}_i$ and $e_j\mathcal{C} = \mathcal{B}_j$, if $j \neq i$. Clearly we have $\mathrm{Rad}(e_i\mathcal{C}) = \mathrm{Rad}(\mathcal{C}) \neq 0$. Now a two sided ideal $\mathcal{J}_i$ of $e_i\mathcal{C}$ minimal among those containing $\mathrm{Rad}(e_i\mathcal{C})$ which is a left ideal of $\mathcal{B}_i$ will clearly suffice as $\mathcal{J}$.

For the rest of the proof we assume that $\mathcal{B}$ is a simple algebra. Let V be a simple left $\mathcal{B}$-module, and let D stand for the algebra of $\mathcal{B}$-endomorphisms of V. By Schur's lemma D is a division algebra over the field F and V is a right D-space. Moreover, we have $\mathcal{B} = \mathrm{End}_D V$ and hence $\mathrm{Rad}(\mathcal{C})V \neq 0$. We define the strictly decreasing chain of D-subspaces $V = V_0 \supset V_1 \supset V_2$ by $V_{i+1} = \mathrm{Rad}(\mathcal{C})V_i$, for $i = 0, 1$. From this chain of subspaces we obtain a decreasing chain of subalgebras $\mathcal{B} = \mathcal{B}_0 \supseteq \mathcal{B}_1 \supseteq \mathcal{B}_2$ by letting

$$\mathcal{B}_i = \{x \in \mathcal{B} \mid xV_j \subseteq V_j \text{ for } j = 0, \ldots, i\}.$$

Here $\mathcal{B} \neq \mathcal{B}_1$ follows from $\mathcal{B} = \mathrm{End}_D V$. Moreover, $\mathcal{B}_2 \supseteq \mathcal{C}$ implies that $\mathcal{B}_1 = \mathcal{B}_2 = \mathcal{C}$. We infer that $V_2 = 0$ and $(\mathrm{Rad}(\mathcal{C}))^2 = 0$.

Then the annihilator $\mathcal{J} = \{x \in \mathcal{B} \mid xV_1 = 0\}$ is properly contained in $\mathcal{B}_1 = \mathcal{C}$, and in fact is a two-sided ideal of $\mathcal{C}$. It is also obvious, that $\mathcal{J}$ is a left ideal of $\mathcal{B}$, and this implies that $\mathcal{J} \supset \mathrm{Rad}(\mathcal{C})$. From $\mathcal{B} = \mathrm{End}_D V$ we obtain that $\mathcal{C}/\mathrm{Rad}(\mathcal{C}) \cong \mathrm{End}_D V_1 \oplus \mathrm{End}_D V/V_1$. Thus, $\mathcal{C}/\mathrm{Rad}(\mathcal{C})$ is a semisimple algebra with exactly two minimal ideals, implying the minimality of $\mathcal{J}$ over $\mathrm{Rad}(\mathcal{C})$. □

Theorem 6.3.5 *Let $\Lambda \subset \Gamma$ be R-orders in $\mathcal{A}$. Suppose that Λ is extremal and Γ is minimal among the R-orders properly containing Λ. Then there exists an ideal $\mathcal{I}$ of Λ minimal among those containing* $\mathrm{Rad}(\Lambda)$ *such that $\mathcal{O}_l(\mathcal{I}) \supseteq \Gamma$.*

Proof By Propositions 6.3.3 and 6.2.3 we have that $\mathcal{C} = \Lambda/\mathrm{Rad}(\Gamma)$ is a maximal proper subalgebra of the semisimple $F = R/P$-algebra $\mathcal{B} = \Gamma/\mathrm{Rad}(\Gamma)$. Moreover $\mathrm{Rad}(\mathcal{C}) \neq 0$, since $\Lambda \subset \Gamma$ and Λ is extremal. We can apply Lemma 6.3.4. There exists a minimal ideal $\mathcal{J}$ of $\mathcal{C}$ above $\mathrm{Rad}(\mathcal{C})$ such that $\mathcal{J}$ is a left ideal in $\mathcal{B}$. Now $\mathcal{I}$, the inverse image of $\mathcal{J}$ with respect to the natural map $\Gamma \to \mathcal{B}$ clearly satisfies the requirements of the theorem. □

6.4 Computing maximal orders

Let $R = \mathbb{Z}$ or $R = \mathbb{F}_q[X]$, K be the field of quotients of R (i.e., $K = \mathbb{Q}$, or $K = \mathbb{F}_q(X)$, respectively), L be a finite separable extension of K, $\mathcal{A}$ be a finite dimensional separable L-algebra, and Λ is an R-order in $\mathcal{A}$. Let D stand for the integral closure of R in L. Suppose that $\mathcal{A}$ is given by structure constants over L and Λ is given by an R-basis.

Suppose further that we are given a prime element $\pi \in R$ (π is a rational prime in case $R = \mathbb{Z}$, or an irreducible polynomial over $\mathbb{F}_q$ in case $R = \mathbb{F}_q[X]$).

Theorem 6.4.1 *There exists an f-algorithm running in time* $(\mathrm{size}(\mathcal{A}) + \mathrm{size}(\Lambda) + \mathrm{size}(\pi))^{O(1)}$ *that produces an R-basis of an R-order $\Gamma \supset \Lambda$ such that $\Gamma_{(\pi)} \supset \Lambda_{(\pi)}$ provided that Λ is not maximal at π. The f-oracle is employed to factor polynomials over $R/(\pi)$.*

Proof We shall test first whether $\Lambda_{(\pi)}$ is an extremal $R_{(\pi)}$-order by checking if $\mathcal{O}_l(\mathrm{Rad}(\Lambda_{(\pi)})) = \Lambda_{(\pi)}$. If not, then we construct an R-order $\Gamma \supset \Lambda$ such that $\Gamma_{(\pi)} \supset \Lambda_{(\pi)}$. If $\Lambda_{(\pi)}$ passes the test, then we use the following test based on Theorem 6.3.5. If there exists an ideal $\mathcal{J}$ minimal among the ideals properly containing $\mathrm{Rad}(\Lambda_{(\pi)})$ such that $\mathcal{O}_l(\mathcal{J}) \supset \Lambda_{(\pi)}$, then we construct an R-order $\Gamma \supset \Lambda$ such that $\Gamma_{(\pi)} \supset \Lambda_{(\pi)}$. Otherwise we correctly conclude that Λ is maximal at π.

As for the first test, we compute the inverse image $\mathcal{I} \subseteq \Lambda$ of $\mathrm{Rad}(\Lambda_{(\pi)})$ with respect to the embedding $\Lambda \to \Lambda_{(\pi)}$. By Lemma 6.1.9, Λ passes the first test if and only if $\mathcal{O}_l(\mathcal{I}) = \Lambda$. Otherwise $\Gamma = \mathcal{O}_l(\mathcal{I})$ is an order containing Λ such that $\Gamma_{(\pi)} \supset \Lambda_{(\pi)}$.

We shall work with the finite algebra $\mathcal{B} = \Lambda/\pi\Lambda$ over the finite field $F = R/(\pi)$. We have $\mathrm{size}(F) = \mathrm{size}(\pi)^{O(1)}$, $\dim_F \mathcal{B} \leq \dim_K \mathcal{A} = n$ and structure constants for $\mathcal{B}$ are easily obtained, $\mathrm{size}(\mathcal{B}) = (\mathrm{size}(\mathcal{A}) + \mathrm{size}(\Lambda) + size(F))^{O(1)}$. From Propositions 6.2.1 and 6.1.5 we infer that $\mathcal{I}$ is the inverse image of $\mathrm{Rad}(\mathcal{B})$ with respect to the canonical map $\Lambda \to \mathcal{B}$. $\mathrm{Rad}(\mathcal{B})$ can be computed in deterministic time $(n + \mathrm{size}(F))^{O(1)}$ with the method of [Ró2]. From an F-basis of $\mathrm{Rad}(\mathcal{B})$ we can efficiently find an R-basis of $\mathcal{I}$. (Note that any R-submodule M such that $\pi\Lambda \subseteq M \subseteq \Lambda$ has a basis of size bounded by $(\mathrm{size}(\Lambda)+\mathrm{size}(\pi))^{O(1)}$). Also, by Proposition 6.1.8 we can compute $\mathcal{O}_l(\mathcal{I})$ efficiently. This finishes the description of the first test.

The second test can be treated in a similar way. Let $\mathcal{J}_1, \ldots, \mathcal{J}_m$ denote the minimal ideals of $\mathcal{B}$ which contain $\mathrm{Rad}(\mathcal{B})$. Note that these ideals are the inverse images, with respect to the canonical map $\phi : \mathcal{B} \to \mathcal{B}/\mathrm{Rad}(\mathcal{B})$, of the minimal ideals of the semisimple algebra $\mathcal{B}/\mathrm{Rad}(\mathcal{B})$. We have $m \leq n$. Let $\mathcal{I}_i$ denote the inverse image in Λ of $\mathcal{J}_i$ with respect to the map $\Lambda \to \mathcal{B}$. Propositions 6.1.5 and 6.2.1 imply that $\mathcal{I}_1, \ldots, \mathcal{I}_m$ are also the inverse images of the minimal ideals of $\Lambda_{(\pi)}$ over $\mathrm{Rad}(\Lambda_{(p)})$. As in the first case, we obtain that we have to compute the rings $\mathcal{O}_l(\mathcal{I}_i)$ for $i = 1, \ldots, m$. We can stop when $\Lambda \subset \mathcal{O}_l(\mathcal{I}_i)$ is detected because then we have an order properly containing Λ.

The ideals $\mathcal{J}_i$ are obtained by the f-algorithm of Friedl and Rónyai (see [FR, Ró2]). The time requirement is $(n + \log p)^{O(1)}$ and we call the f-oracle to factor polynomials over $\mathbb{F}_p$. From the ideals $\mathcal{J}_i$ the ideals $\mathcal{I}_i$ and the rings $\mathcal{O}_l(\mathcal{I}_i)$ can be computed in deterministic time $(\mathrm{size}(\mathcal{A}) + \mathrm{size}(\Lambda) + \log p)^{O(1)}$. The description of the algorithm and the proof of Theorem 6.4.1 is complete. □

With the method of Theorem 6.4.1 we can construct a maximal R-order in $\mathcal{A}$ as follows. First we need a starting R-order. Let $a_1, \ldots, a_n$ be the input basis of $\mathcal{A}$ over K. Let d be lowest common denominator of the structure constants with respect to this basis. Then the R-module Λ generated by $1_{\mathcal{A}}, da_1, \ldots, da_n$ is an R-order in $\mathcal{A}$. (This is the same trick that was used in Section 1.2 to make structure constants integral.) We put

$u = \det\left(\mathrm{tr}_{\mathcal{A}/K}(d^2 a_i a_j)\right)_{i,j=1}^n$. The elements $\mathrm{tr}_{\mathcal{A}/K}(a_i a_j)$ can be computed if we know the Wedderburn decomposition of $\mathcal{A}$ over K. In the case $K = \mathbb{Q}$, the Wedderburn decomposition can be computed in deterministic polynomial time (cf. [FR]). In the case $K = \mathbb{F}_q(X)$, this task can be performed by the deterministic polynomial time f-algorithm of Theorem 4.1.3. Now u is a multiple of $\mathrm{disc}(\Lambda)$, whence, by Proposition 6.1.3, Λ is maximal at every prime π not dividing u. Let S be the set of primes in R dividing u. S is obtained by factoring u in R.

Repeated application of Theorem 6.4.1 gives a sequence of R-orders

$$\Lambda = \Gamma_0 \subset \Gamma_1 \subset \ldots \subset \Gamma_m$$

until a maximal R-order is obtained. By Proposition 6.1.10, $m \leq \frac{1}{2}\log_2|u|$ if $K = \mathbb{Q}$ and $m \leq \frac{1}{2}\deg u$ if $K = \mathbb{F}_q[X]$. We can control sizes during the iteration. By Proposition 6.1.3 we have $\Lambda \subseteq \Gamma_j \subseteq \frac{1}{u}\Lambda$, therefore Γ_j can be represented by an R-basis admitting a short description.

Theorem 6.4.2 *Let $\mathcal{A}$ be a finite dimensional separable algebra over K given by structure constants.*

(a) If $K = \mathbb{Q}$, then a maximal $\mathbb{Z}$-order Λ can be constructed by an ff-algorithm running in time $\mathrm{size}(\mathcal{A})^{O(1)}$.

(b) If $K = \mathbb{F}_q(X)$, then a maximal $\mathbb{F}_q[X]$-order Λ can be constructed by an f-algorithm running in time $\mathrm{size}(\mathcal{A})^{O(1)}$. □

Corollary 6.4.3 *Let L be a finite separable extension of K and $\mathcal{A}$ a be finite dimensional central simple algebra over L. Let D denote the integral closure of R in L. Suppose that $\mathcal{A}$ is given by structure constants over L. Then a maximal D-order Λ in $\mathcal{A}$ can be constructed*

(a) by an ff-algorithm running in time $\mathrm{size}(\mathcal{A})^{O(1)}$, *if $R = \mathbb{Z}$.*

(b) by an f-algorithm running in time $\mathrm{size}(\mathcal{A})^{O(1)}$, *if $R = \mathbb{F}_q[X]$.*

Proof From K and the structure constants of $\mathcal{A}$ over L we can readily obtain structure constants of $\mathcal{A}$ over K. With the method of Theorem 6.4.2 we compute an R-basis of a maximal R-order Λ of $\mathcal{A}$. By Lemma 6.1.6 we conclude that Λ is a maximal D-order as well. □

Corollary 6.4.3 (a) gives an affirmative answer to the question proposed in [Ró3].

6.5 Computing indices

First we recall some standard material related to valuations and completions (cf. [Re], Section 5). Assume that $\mathcal{A}$ is a central simple K-algebra of dimension n^2. Let ϕ be a valuation of K. We consider the completions K_ϕ and $\mathcal{A}_\phi = K_\phi \otimes_K \mathcal{A}$, respectively. It is easy to see that $\mathcal{A}_\phi$ is a central simple K_ϕ-algebra of dimension n^2, therefore $\mathrm{index}(\mathcal{A}_\phi) = m_\phi$ for some $m_\phi | n$, i.e., $\mathcal{A}_\phi \cong \mathrm{M}_{n/m_\phi}(\mathcal{D}_\phi)$, where $\mathcal{D}_\phi$ is a central skewfield over K_ϕ with $\dim_K \mathcal{D}_\phi = m_\phi^2$. We call the index m_ϕ the *local index* of $\mathcal{A}$ at the valuation ϕ. Since for every central simple K-algebra $\mathcal{B}$ we have $\mathrm{M}_r(\mathcal{B})_\phi = \mathrm{M}_r(\mathcal{B}_\phi)$, the local index m_ϕ is in fact a divisor of $\mathrm{index}(\mathcal{A})$.

Assume that ϕ is a nonarchimedean valuation of K. Then $R_\phi = \{x \in K^* | \phi(x) \leq 1\}$ is a subring of K, called the valuation ring of ϕ. $P_\phi = \{x \in K^* | \phi(x) < 1\}$ is the unique maximal ideal in R_ϕ, called the valuation ideal. If ϕ is a discrete valuation, then R_ϕ is a discrete valuation ring, i.e., the only prime ideal of R_ϕ is P_ϕ. Assume further that the residue class field R_ϕ/P_ϕ is finite. This holds for every nonarchimedean valuation of a global field K. (In fact, global fields can be characterized with this property.) The following statement, based on the classification of division algebras over local fields (cf. [Re], Chapter 3) and the theory of maximal orders over discrete valuation rings (cf. [Re], Chapter 5), relates the local index m_ϕ of $\mathcal{A}$ to the structure of maximal R_ϕ-orders in $\mathcal{A}$.

Proposition 6.5.1 *Let ϕ be a discrete valuation of K such that the residue class field R_ϕ/P_ϕ is finite and Λ be a maximal R_ϕ-order in the central simple K-algebra $\mathcal{A}$ of dimension n^2. Then* $\mathrm{Rad}(\Lambda)$ *is the a unique maximal two-sided ideal in* Λ,

$$P_\phi \Lambda = (\mathrm{Rad}(\Lambda))^{m_\phi}, \quad \textit{and} \quad \Lambda/\mathrm{Rad}(\Lambda) \cong \mathrm{M}_t(\mathcal{B}),$$

where $\mathcal{B}$ is a field extension of R_ϕ/P_ϕ of degree m_ϕ and $t = n/m_\phi$.

Proof Let S_ϕ be the valuation ring of the valuation ϕ of the field K_ϕ, and Q_ϕ be the maximal ideal of S_ϕ. Let $\Omega = S_\phi \otimes_{R_\phi} \Lambda$. By [Re], Thm. 11.5, Ω is a maximal S_ϕ-order in $\mathcal{A}_\phi$, and by [Re], Thm. 18.7, $\mathrm{Rad}(\Lambda) = \Lambda \cap \mathrm{Rad}(\Omega)$, whence $\Lambda/\mathrm{Rad}(\Lambda) \cong \Omega/\mathrm{Rad}(\Omega)$.

By [Re], Thm. 17.3, Ω is conjugate by an inner automorphisms of $\mathcal{A}_\phi$ to the order $\mathrm{M}_t(\Delta)$, where $\mathcal{A}_\phi \cong \mathrm{M}_t(\mathcal{D})$, $\mathcal{D}$ is a central skewfield over $\mathcal{A}_\phi$ of index m_ϕ and Δ is the unique maximal S_ϕ-order in $\mathcal{D}$. We have

$$\Lambda/\mathrm{Rad}(\Lambda) \cong \Omega/\mathrm{Rad}(\Omega) \cong \mathrm{M}_t(\Delta/\mathrm{Rad}(\Delta)).$$

By [Re], Thm 14.3, $\mathcal{B} = \Delta/\mathrm{Rad}(\Delta)$ is an extension field of degree m_ϕ of $S_\phi/Q_\phi \cong R_\phi/P_\phi$, and $\mathrm{Rad}(\Delta)^{m_\phi} = Q_\phi \Delta$. If we identify Ω with $\mathrm{M}_t(\Delta)$, we have $\mathrm{Rad}(\Omega) = \mathrm{M}_t(\mathrm{Rad}(\Delta))$ and

$$Q_\phi \Omega = \mathrm{M}_t(Q_\phi \Delta) = \mathrm{M}_t((\mathrm{Rad}(\Delta))^{m_\phi}) = \mathrm{M}_t((\mathrm{Rad}(\Delta)))^{m_\phi} = (\mathrm{Rad}(\Omega))^{m_\phi}.$$

It follows that

$$P_\phi\Lambda = \Lambda \cap Q_\phi\Omega = \Lambda \cap (\mathrm{Rad}(\Omega))^{m_\phi}) = (\Lambda \cap \mathrm{Rad}(\Omega))^{m_\phi} = (\mathrm{Rad}(\Lambda))^{m_\phi}.$$

□

Assume that R is a Dedekind domain with quotient field K. If R is contained in the valuation ring R_ϕ of ϕ (i.e., $\phi(x) \leq 1$ for every $x \in R$) then $P = R \cap P_\phi = \{x \in R | \phi(x) < 1\}$ is a maximal ideal in R and ϕ is equivalent to the usual P-adic valuation ϕ_P of K. We say in this case that ϕ corresponds to the prime ideal P of R. We call the local index m_{ϕ_P} the local index at P and denote it by m_P.

From now on, we assume that $R = \mathbb{Z}$ or $R = \mathbb{F}_q[X]$, K is the quotient field of R and $\mathcal{A}$ is a separable K-algebra. First we give a method to compute the local indices of simple separable algebras over global fields. (Recall that the simple components of a semisimple $\mathbb{Q}$-algebra can be computed by the deterministic polynomial time method of [FR], while the analogous task for algebras over $\mathbb{F}_q(x)$ can be done by the deterministic polynomial time f-algorithm of Theorem 4.1.3.)

Assume that the center of $\mathcal{A}$ is L, a finite separable extension of K. Let D be the integral closure of R in L. The following simple statement tells us, that we do not have to care about local indices at primes not dividing the discriminant of an order.

Proposition 6.5.2 *Let $\mathcal{A}$ be a central simple algebra over L, a finite separable extension of K. Let D be the integral closure of R in $\mathcal{A}$. Assume that Λ is an R-order in $\mathcal{A}$. Assume that for some prime ideal P in D above the prime $\pi \in R$ we have $m_P > 1$. Then $\pi | \mathrm{disc}_K(\Lambda)$.*

Proof Assume that π does not divide the discriminant. The $R/(\pi)$-algebra $\Lambda/\pi\Lambda$ has a nonsingular bilinear trace form, whence it is semisimple. Therefore its factor $\Lambda/P\Lambda \cong \Lambda_P/P\Lambda_P$ is also semisimple implying that $\mathrm{Rad}(\Lambda)_P = P\Lambda_P$. On the other hand, by Proposition 6.5.1 we have $P\Lambda_P = (\mathrm{Rad}(\Lambda)_P)^{m_P}$, giving that $m_P = 1$. □

Note that in [Re], Theorem 32.1 an exact formula is given for the factorization of the discriminant of a maximal order. This would make possible to compute local indices from the discriminant of a maximal order, as it is done in [Ró3]. We show instead a method based on the structure of the factor of a locally maximal order by the radical.

Proposition 6.5.3 *Given a finite separable extension L of K, a central simple algebra $\mathcal{A}$ over L, and a prime $\pi \in R$, we can compute the set of local indices of $\mathcal{A}$ at primes of D above π by a deterministic polynomial time f-algorithm.*

Proof With the f-algorithm of Theorem 6.4.1, we first compute an R-order Λ in $\mathcal{A}$ which is locally maximal at π. Hereon we perform computations in the factor ring $\mathcal{B} = \Lambda/\pi\Lambda$. $\mathcal{B}$ is a finite dimensional algebra over the finite field $R/(\pi)$. We compute the subalgebra $\mathcal{C} \cong (\Lambda \cap L)/(\pi\Lambda \cap L)$, the image of $\Lambda \cap L$ by the natural map. We use the method of [Ró2] to compute $\mathrm{Rad}(\mathcal{C})$. Using the f-algorithm of [FR] for decomposing $\mathcal{C}/\mathrm{Rad}(\mathcal{C})$, we can find the maximal ideals of $\mathcal{C}$. By Lemma 6.1.7, these ideals correspond to the prime ideals of D over π. For every such ideal M (corresponding to the D-ideal P), we compute the factor ring $\mathcal{B}/M\mathcal{B}$. By Lemma 6.1.7, this ring is isomorphic to $\Gamma/P\Gamma$, where $\Gamma = D\Lambda$. By Proposition 6.5.1, the radical-free part of $\Gamma_P/P\Gamma_P \cong \Gamma/P\Gamma$ is a full matrix algebra over an extension of degree m_P of D/P. Therefore, by Lemma 6.1.7, the index m_P can be obtained as the dimension of the center of the radical-free part of $\mathcal{B}/M\mathcal{B}$ over the field $\mathcal{C}/M$. $\square$

The last two statements suggest a method to compute the set of all local indices for valuations corresponding to prime ideals in D based on factoring the discriminant of a starting order. This leads to a deterministic polynomial time ff-algorithm in case $R = \mathbb{Z}$, or a deterministic polynomial time f-algorithm in case $R = \mathbb{F}_q[X]$. Note that we do not need to compute the ring D.

We show how to compute the local indices for valuations not corresponding to primes in R. If ϕ is archimedean, then, at least in the cases we considered, we may assume that $K = \mathbb{Q}$, and the completion L_ϕ corresponds to an embedding $L \to \mathbb{C}$. In that case $L_\phi \cong \mathbb{C}$ or $L_\phi \cong \mathbb{R}$, and the only possible proper skewfield is that of the Hamiltonian quaternions (only in the case $L_\phi \cong \mathbb{R}$). We can use Eberly's randomized polynomial method ([Eb3]) to determine the set of local indices at nonarchimedean primes. Since in [Eb3] randomization is used only at the point of finding *splitting elements*, using the results of Section 5.5 (or even using [Ró5]) we can derandomize it. The nonarchimedean valuations of an algebraic number field L correspond to prime ideals in the ring D of algebraic integers in L, therefore they can be treated by the method of Proposition 6.5.1.

If $K = \mathbb{F}_q(x)$, then every valuation of L is discrete. If $\phi(X) \leq 1$ then $\mathbb{F}_q[X]$ and its integral closure D belong to the valuation ring, the case treated by the method of Proposition 6.5.1. On the other hand, $\phi(X) > 1$ implies

$$\phi(\frac{g(X)}{h(X)}) = 2^{(\deg g(X) - \deg h(X)) \log_2 \phi(X)}.$$

It follows that $\phi(\frac{1}{X}) < 1$ and ϕ corresponds to the prime ideal $(\frac{1}{X})$ in the ring $\mathbb{F}_q[\frac{1}{X}]$. This case can be treated as above by using $\mathbb{F}_q[\frac{1}{X}]$ in place of $\mathbb{F}_q[X]$.

We use the following reformulation (cf. [Re], Theorem 32.19) of the celebrated and deep Albert–Brauer–Hasse–Noether theorem.

Theorem. *If $\mathcal{A}$ is a central simple algebra over the global field L, then*

$$\text{index}(\mathcal{A}) = \text{lcm}_\phi m_\phi,$$

where ϕ runs over the valuations of L.

Corollary 6.5.4 *The index of the separable simple algebra $\mathcal{A}$ over K can be computed by a polynomial time*

(a) ff-algorithm if $K = \mathbb{Q}$;

(b) f-algorithm if $K = \mathbb{F}_q[X]$.

Next we give an application to group representations. Suppose that K is an algebraic number field. Let $G \leq GL_n(K)$ be a finite matrix group given by a finite set $\Gamma \subseteq GL_n(K)$ of matrices generating G. We assume the dense representation (cf. Section 1.1) of K, elements of K and matrices over K. Note that finiteness of G can be tested by the deterministic polynomial time method of [BBR]. By Maschke's theorem, the action of G on $V = K^n$, the K-space of column vectors of length n over K, is completely reducible. We are interested in computing the dimensions of the irreducible constituents.

Theorem 6.5.5 *We can compute the dimensions of the irreducible constituents of the action of G on K^n in deterministic polynomial time.*

Proof Let $\mathcal{A}$ be the K-subalgebra of $\mathrm{M}_n(K)$ generated by Γ. A basis of $\mathcal{A}$ can be computed in deterministic polynomial time. By Maschke's theorem, $\mathcal{A}$ is semisimple. With the deterministic polynomial time method of [FR] we can compute the simple components $\mathcal{A}_1, \ldots, \mathcal{A}_r$ of $\mathcal{A}$ together with the centers $K_1, \ldots, K_r$ of these components. Setting $V_1 = \mathcal{A}_1 V, \ldots, V_r = \mathcal{A}_r V$, we have $V = V_1 \oplus \ldots \oplus V_r$, where for every $i = 1, \ldots, r$, the irreducible constituents of V_i are isomorphic to a minimal left ideal L_i of $\mathcal{A}_i$, whence the dimension of an irreducible constituent of V_i is

$$\dim_K L_i = \dim_K K_i \dim_{K_i} L_i = \dim_K K_i \text{index}\mathcal{A}_i \sqrt{\dim_{K_i} \mathcal{A}_i}.$$

Applying the projections $\mathcal{A} \to \mathcal{A}_i$ to the generators in Γ, our task is reduced to r instances of the problem of computing the index of $\mathcal{A}$ in the special case where the K-algebra $\mathcal{A}$ generated by Γ is a central simple algebra over K (this K may be different from the field we started with). In that case $\mathcal{A}$ is isomorphic to a simple component of the group algebra KG. Let D stand for the ring of algebraic integers in K. It is known (cf. [Re], Section 41) that if p is a prime not dividing $|G|$, and P is a prime in D above p, then the local index of every simple component of KG at P is 1. In fact, the $\mathbb{Z}$-lattice Λ generated by D and Γ is a D-order in $\mathcal{A}$ with $\Lambda/P\Lambda \cong \mathrm{M}_{\sqrt{\dim_K \mathcal{A}}}(D/P)$.

It follows that it is enough to compute the local indices of $\mathcal{A}$ at the archimedean valuations and at the primes P of D above rational primes p such that p divides $|G|$. It is known that these primes are small. Indeed, using the regular representation of K over $\mathbb{Q}$, we may think of elements of G as elements of $\mathrm{M}_{n\dim_{\mathbb{Q}} K}(\mathbb{Q})$. Let $g \in G$ be an element of G of prime order p. Then g is a root of the cyclotomic polynomial $\frac{X^p-1}{X-1}$, which is irreducible over $\mathbb{Q}$, therefore $\frac{X^p-1}{X-1}$ divides the characteristic polynomial of g, whence $p-1 \le n\dim_{\mathbb{Q}} K$. Therefore it suffices to compute the local indices at primes above rational primes $p \le n\dim_{\mathbb{Q}} K + 1$. Using the f-algorithm of 6.4.1, for every prime $p \le n\dim_{\mathbb{Q}} K + 1$ we compute an order Λ in $\mathcal{A}$ such that Λ is locally maximal at p. The f-oracle for factoring polynomials over $\mathbb{F}_p$ can be replaced with the deterministic $(\dim_{\mathbb{Q}} \mathcal{A} + p)^{O(1)}$-time method of Berlekamp (cf. [Ber3]). $\square$

Applying this to the regular representation of G we immediately obtain

Corollary 6.5.6 *Let K be an algebraic number field given by the monic minimal polynomial $f \in \mathbb{Z}[x]$ of an integral element $\alpha \in K$ such that $K = \mathbb{Q}(\alpha)$. Let G be a finite group given by its multiplication table. Then the degrees of the irreducible representations of G over K can be computed in deterministic time* $(\mathrm{size}(f) + |G|)^{O(1)}$. $\square$

Chapter 7

Finding isomorphism with $\mathrm{M}_2(\mathbb{Z})$

The problem of finding a minimal left ideal in a simple algebra remains open in general; even it is not known if a zero divisor of polynomial size exists. Rónyai [Ró1] has proved that the simplest case, i.e., finding a zero divisor in a simple non-commutative 4-dimensional $\mathbb{Q}$-algebra is essentially as hard as the problem of decomposing integers into prime factors. The results of the previous chapter suggest that finding certain maximal orders may bring us closer to finding minimal left ideals. Here, based on the paper [ISz] (joint work with Ágnes Szántó), we present evidence in favour of this approach: If we are given a maximal order in an algebra isomorphic to $\mathrm{M}_2(\mathbb{Q})$, i.e., a subring isomorphic to $\mathrm{M}_2(\mathbb{Z})$, then we can find a zero divisor in it in deterministic polynomial time.

Our method is based on a kind of basis reduction with respect to the bilinear trace form of the algebra. The algorithm may be of independent interest. A. K. Lenstra, H. W. Lenstra and L. Lovász (see [LLL] and [Lo]) have constructed a polynomial time algorithm to reduce a positive definite integral quadratic form, i.e., to produce a basis of a full lattice L in an n-dimensional $\mathbb{R}$-vector space V equipped with a positive definite bilinear function with integer values on L, such that the matrix of the bilinear function with respect to this basis consists of integers of absolute values bounded by a constant depending only on the discriminant and the dimension. It turns out, that their algorithm can be extended to the case of non-degenerate indefinite forms. The original algorithm has to be modified only at points where isotropic vectors occur.

Section 7.1 contains the definitions, the reduction algorithm and the main properties of reduced bases. In Section 7.2, we show how a basis reduced with respect to the bilinear trace form of a ring isomorphic to $\mathrm{M}_2(\mathbb{Z})$ can be used to find a zero divisor.

7.1 Basis reduction

A *lattice* L in $\mathbb{R}^n$ is a free abelian subgroup generated by an $\mathbb{R}$-basis of $\mathbb{R}^n$. A *basis* of L is a free generating set, i.e., an $\mathbb{R}$-basis of $\mathbb{R}^n$ generating the additive group L. Let L be a lattice in $V = \mathbb{R}^n$ and $\langle\ ,\ \rangle : V \times V \to \mathbb{R}$ a symmetric bilinear function taking integer values on L. The *discriminant* of L is defined by

$$\mathrm{disc}(L) := |\det\left(\langle b_i, b_j\rangle\right)_{i,j=1}^n|$$

where $(b_1, \ldots, b_n)$ is a basis of L. It is known that the discriminant is independent of the choice of the basis of L, and $\mathrm{disc}(L) \in \mathbb{N}$. The function (form) $\langle\ ,\ \rangle$ is called *degenerate* if the discriminant $\mathrm{disc}(L)$ is zero. In this paper we assume that $\langle\ ,\ \rangle$ is *non-degenerate*, i.e., $\mathrm{disc}(L) \neq 0$. A nonzero vector $x \in V$ is called *isotropic* if $\langle x, x\rangle = 0$, anisotropic if $\langle x, x\rangle \neq 0$. Since $\langle\ ,\ \rangle$ is assumed to be non-degenerate, isotropic vectors can exist only in the *indefinite* case, i.e. when $\langle x, x\rangle$ takes positive as well as negative values.

Our aim is to find a basis $(b_1, \ldots, b_n)$ of L such that the matrix $(\langle b_i, b_j\rangle)_{i,j=1}^n$ consists of integers of "small" absolute values. There exist classical reduction methods (introduced by Hermite for definite forms; for an extension to indefinite forms the reader is referred to [Ca], Section 9.3) based on choosing b_1 such that $|\langle b_1, b_1\rangle|$ is as small as possible and then recursively continuing the procedure in the component of L orthogonal to b_1. The main difficulty with this classical approach is in finding a shortest vector. This task is known to be NP-hard. The key idea in [LLL] is to replace the notion of reducedness in the sense of Hermite by an algorithmically tractable requirement.

Let us try first to imitate the LLL reduction algorithm and to extend the main results. To ensure correctness, we assume that all the denominators appearing are nonzero and postpone the discussion of the modifications necessary to cover the exceptional cases.

We will extensively use some notational conventions related to the *Gram–Schmidt orthogonalization* of a basis $(b_1, \ldots, b_n)$ of L.

$$b_1^* := b_1, \quad b_i^* := b_i - \sum_{j=1}^{i-1} \frac{\langle b_i, b_j^*\rangle}{\langle b_j^*, b_j^*\rangle} b_j^* \quad (i = 2, \ldots, n),$$

i.e., b_i^* is the component of b_i orthogonal to the subspace generated by $\{b_1, \ldots, b_{i-1}\}$. For the quotients $\frac{\langle b_i, b_j^*\rangle}{\langle b_j^*, b_j^*\rangle}$, $1 \leq j \leq i \leq n$ we will also use the notation μ_{ij}. We have $b_i = \sum_{j=1}^{i} \mu_{ij} b_j^*$. More generally, let $b_i(j)$ denote the component of b_i orthogonal to the subspace generated by $b_1, \ldots, b_{j-1}$, i.e.

$$b_i(j) = \begin{cases} b_i - \sum_{t=1}^{j-1} \frac{\langle b_i, b_t^*\rangle}{\langle b_t^*, b_t^*\rangle} b_t^* & \text{if } j \leq i \\ 0 & \text{if } j > i. \end{cases}$$

In particular, $b_i^* = b_i(i)$. Note that the vector b_i^* depends only on b_i and the sublattice L_{i-1} generated by $b_1, \dots, b_{i-1}$.

A basis is $b_1, \dots, b_n$ is called *reduced* in the sense of LLL if

(1) $\quad |\mu_{ij}| \le \frac{1}{2}$ $\qquad$ for every $1 \le j < i \le n$, and

(2) $\quad |\langle b_i(i), b_i(i)\rangle| \le \frac{4}{3}|\langle b_{i+1}(i), b_{i+1}(i)\rangle|$ $\qquad$ for every $1 \le i < n$.

Bases satisfying property (1) are called *proper*. Besides properness, the following consequence of (1) and (2) is crucial in estimates:

(2′) $\quad |\langle b_{i+1}^*, b_{i+1}^*\rangle| \ge c|\langle b_i^*, b_i^*\rangle|$ $\qquad$ for every $1 \le i < n$

with some constant $c < 1$. (Actually, $c = \frac{1}{2}$.) The proof of this is identical to the proof for definite forms in [LLL]. In fact, (2′) follows from

$$\begin{aligned} \frac{3}{4}|\langle b_i^*, b_i^*\rangle| &\le |\langle b_{i+1}(i), b_{i+1}(i)\rangle| = |\langle b_{i+1}^*, b_{i+1}^*\rangle + \mu_{i+1,i}^2\langle b_i^*, b_i^*\rangle| \\ &\le |\langle b_{i+1}^*, b_{i+1}^*\rangle| + \frac{1}{4}|\langle b_i^*, b_i^*\rangle|. \end{aligned}$$

From a basis $b_1, \dots, b_n$, it is easy to construct a proper basis such that the Gram–Schmidt orthogonalization results in the same basis $b_1^*, \dots, b_n^*$:

```
for i := 1 to n do
    for j := i − 1 downto 1 do
        b_i := b_i − m b_j where m is the nearest integer to μ_ij
```

A reduced basis is obtained by repeating the following rounds:
- properness is achieved/maintained with the above procedure;
- we swap b_i and b_{i+1} if property (2) does not hold for i.

The crucial part of the analysis of the algorithm relies on the fact that the quantity

$$\prod_{i=1}^{n} \mathrm{disc}(L_i) = \prod_{i=1}^{n}\prod_{j=1}^{i} |\langle b_j^*, b_j^*\rangle|$$

is reduced by a factor less than $\frac{3}{4}$ in every execution of the swapping step.

Now we have to treat the "exceptions": First of all, the definition of b_i^* and $b_i(j)$ will only work as long as we have $\langle b_i^*, b_i^*\rangle \ne 0$, since the orthogonalization procedure stops as soon as an isotropic vector b_i^* is found. This happens at the first index i such that our form restricted to the subspace L_i is degenerate, since $\mathrm{disc}(L_i) = \prod_{j=1}^{i}\langle b_j^*, b_j^*\rangle$. A basis $(b_1, \ldots, b_n)$ is called *singular*, if there is an i, $1 \leq i \leq n$, such that $\mathrm{disc}(L_i) = 0$. We intend to work only with *nonsingular* bases, i.e., when $\langle\ ,\ \rangle$ restricted to L_i is non-degenerate for $i = 1, \ldots, n$.

Finding a nonsingular basis very close to a given singular one appears to be an easy task: Let i be the lowest index such that $\langle b_i^*, b_i^*\rangle = 0$. The vectors $b_1^*, \ldots, b_i^*$ together with $b_{i+1}(i), \ldots, b_n(i)$ form a basis of V. The non-degeneracy of $\langle\ ,\ \rangle$ implies that b_i^* cannot be orthogonal to every element of that basis, hence $i < n$ and there exists some j, $i < j \leq n$, such that $\langle b_i^*, b_j(i)\rangle \neq 0$. Since

$$\langle b_i^*, b_j(i)\rangle = \frac{1}{4}(\langle b_i^* + b_j(i), b_i^* + b_j(i)\rangle + \langle b_i^* - b_j(i), b_i^* - b_j(i)\rangle),$$

there exists $\epsilon = \pm 1$ such that $b_i^* + \epsilon b_j(i)$ is anisotropic. But this vector is nothing else than the component of the lattice vector $b_i + \epsilon b_j$ orthogonal to the subspace generated by $b_1, \ldots, b_{i-1}$, thus by replacing b_i with $b_i + \epsilon b_j$, we obtain a basis of L such that for all $1 \leq k \leq i$, $\langle b_k^*, b_k^*\rangle \neq 0$. Iterating this step at most $n-1$ times a nonsingular basis is obtained.

We prepend this procedure to the usual reduction algorithm to ensure that it works with nonsingular bases. Observe that the procedure for making bases proper preserves nonsingularity.

Let us examine the reduction step related to condition (2) of reducedness. Swapping b_i and b_{i+1} does not affect the sublattices L_j for $j \neq i$ and replaces b_i^* with $b_{i+1}(i)$, thus it preserves nonsingularity except when $b_{i+1}(i)$ is isotropic. Then we shall find another lattice vector on the line $b_i - zb_{i+1}$ with short complement $b_i^* - zb_{i+1}(i)$ to replace b_i, except in the (very explicit) case when

$$|\mu_{i+1,i}| = |\frac{\langle b_i^*, b_{i+1}(i)\rangle}{\langle b_i^*, b_i^*\rangle}| = \frac{1}{2}.$$

If $|\mu_{i+1,i}| < \frac{1}{2}$, then we choose z to be the nearest integer to $\frac{1}{2\mu_{i+1,i}}$, so that

$$|z - \frac{1}{2\mu_{i+1}, i}| < \frac{1}{2} < \frac{1}{4|\mu_{i+1,i}|}.$$

This implies that $0 < 1 - 2z\mu_{i+1,i} < \frac{1}{2}$, hence for the complement $b_i^* - zb_{i+1}(i)$ of $b_i - zb_{i+1}(i)$ orthogonal to $b_1, \ldots, b_{i-1}$ we have

$$\frac{|\langle b_i^* - zb_{i+1}(i), b_i^* - zb_{i+1}(i)\rangle|}{|\langle b_i^*, b_i^*\rangle|} = \frac{|\langle b_i^*, b_i^*\rangle - 2z\langle b_i^*, b_{i+1}(i)\rangle|}{|\langle b_i^*, b_i^*\rangle|} = |1 - 2z\mu_{i+1,i}| < \frac{1}{2}.$$

The argument also shows that $b_i^* - zb_{i+1}(i)$ is anisotropic. Thus, the quantity $\prod_{i=1}^n \operatorname{disc}(L_i)$ is reduced here by a factor less than $\frac{1}{2}$, rather than $\frac{3}{4}$. On the other hand, if $|\mu_{i+1,i}| = \frac{1}{2}$, then we have

$$\begin{aligned}
|\langle b_{i+1}^*, b_{i+1}^* \rangle| &= |\langle b_{i+1}(i) - \mu_{i+1,i} b_i^*, b_{i+1}(i) - \mu_{i+1,i} b_i^* \rangle| \\
&= |\mu_{i+1,i}^2 \langle b_i^*, b_i^* \rangle - 2\mu_{i+1,i} \langle b_i^*, b_{i+1}(i) \rangle| \\
&= | - \mu_{i+1,i}^2 \langle b_i^*, b_i^* \rangle| = \frac{1}{4} |\langle b_i^*, b_i^* \rangle|,
\end{aligned}$$

thus $(2')$ is satisfied with $c = \frac{1}{4}$. As a conclusion, the reducedness of a basis $b_1, \ldots, b_n$ with respect to a possibly indefinite form $\langle\ ,\ \rangle$ coincides with the original notion with the following modifications:

(0) the basis must be *nonsingular*,
(1) coincides with the original notion of properness,
(2a) same as (2) *for indices* i, *such that* $\langle b_{i+1}(i), b_{i+1}(i) \rangle \neq 0$,
(2b) $|\mu_{i+1,i}| = \frac{1}{2}$ *for indices* i, *such that* $\langle b_{i+1}(i), b_{i+1}(i) \rangle = 0$.

We have the following method to find a reduced basis:
- *we find a nonsingular basis* and repeat the following rounds:
- achieve/maintain properness
- make changes described above if the test (2a) or (2b) of reducedness fails for some index i.

Theorem 7.1.1 *Given a basis* $(a_1, \ldots, a_n)$ *of a lattice* L *together with the symmetric regular integer matrix* $(\langle a_i, a_j \rangle)_{i,j=1}^n$ *specifying the non-degenerate integral symmetric bilinear function* $\langle\ ,\ \rangle$ *on* L*, the algorithm above finds a reduced basis of* L *in time polynomial in the input size.*

The proof follows very closely the argument from [LLL]. There are only a few points where we have to take into consideration the possible indefiniteness of the form. For this reason, we only sketch here the main differences from the reasoning in [LLL]. The complete proof is relegated to the Appendix.

Proof (*Sketch*) We have seen how to achieve nonsingularity efficienlty. in the case of failure at test (2b), the "length" of the new b_i^* is less than half of the original one, thus the quantity $\prod_{j=1}^n \operatorname{disc}(L_j)$ is reduced by a factor less than $\frac{1}{2}$ (rather than $\frac{3}{4}$). It follows that on the number of arithmetical operations we have bounds similar to the LLL algorithm.

Except for the coordinates of vectors, we obtain similar bounds (they are slightly worse because of indefiniteness) on the numbers (e.g., the values $\langle b_i, b_j \rangle$, numerators and denominators of μ_{ij}) as in the definite case. While definiteness would automatically imply bounds on the coordinates of vectors based purely on bounds on the length, this is no longer true in the indefinite case. Instead, we have to argue as follows: A *single round* of the algorithm results in a transformation matrix with entries of size bounded polynomially by the size of the values $\langle b_i, b_j\rangle$ for the basis $b_1, \ldots, b_n$ that we have at the entry point of the round. But these bounds can be turned into (also polynomial) bounds in terms of the *initial input.* Since we have polynomially many rounds, the result (and the intermediate transformation matrices) will be a product of polynomially many matrices of polynomial size, again a matrix of polynomial size. □

Reduced bases have the following properties:

Theorem 7.1.2 *Let $(b_1, \ldots, b_n)$ be a reduced basis of L. Then for every $1 \le i \le j \le n$,*

$$|\langle b_i, b_j\rangle| \le (2^{n-2} + 2^{n-i-1})\mathrm{disc}(L)^{1/(n-i+1)} \le 2^n \mathrm{disc}(L)^{1/(n-i+1)}.$$

Proof

$$\begin{aligned}
\mathrm{disc}(L) &= \prod_{j=1}^{n} |\langle b_j^*, b_j^*\rangle| = \prod_{j=1}^{i-1} |\langle b_j^*, b_j^*\rangle| \prod_{j=i}^{n} |\langle b_j^*, b_j^*\rangle| = \mathrm{disc}(L_{i-1}) \prod_{j=i}^{n} |\langle b_j^*, b_j^*\rangle| \\
&\ge \prod_{j=i}^{n} |\langle b_j^*, b_j^*\rangle| \ge \prod_{j=i}^{n} 4^{-(j-i)} |\langle b_i^*, b_i^*\rangle|,
\end{aligned}$$

the latter inequality follows from the fact that $(2')$ holds with $c = \frac{1}{4}$. Hence

$$|\langle b_i^*, b_i^*\rangle| \le 2^{n-i}\mathrm{disc}(L)^{1/(n-i+1)}.$$

For $i \le j$ we have

$$\begin{aligned}
|\langle b_i, b_j\rangle| &= |\sum_{k=1}^{i} \mu_{ik}\mu_{jk}\langle b_k^*, b_k^*\rangle| \le \sum_{k=1}^{i-1} \frac{1}{4}|\langle b_k^*, b_k^*\rangle| + |\langle b_i^*, b_i^*\rangle| \\
&\le \frac{1}{4}\sum_{k=1}^{i-1} 2^{n-k}\mathrm{disc}(L)^{1/(n-k+1)} \; + \; 2^{n-i}\mathrm{disc}(L)^{1/(n-i+1)} \\
&\le (\frac{1}{4}\sum_{k=1}^{i-1} 2^{n-k} \; + \; 2^{n-i})\mathrm{disc}(L)^{1/(n-i+1)} \\
&= (2^{n-2} + 2^{n-i} - 2^{n-i-1})\mathrm{disc}(L)^{1/(n-i+1)}.
\end{aligned}$$

□

7.2 Finding zero divisors in $\mathrm{M}_2(\mathbb{Z})$

The smallest examples of non-commutative simple algebras over $\mathbb{Q}$ are of dimension 4. Such an algebra is either a division algebra or isomorphic to $M_2(\mathbb{Q})$, the full ring of 2×2 rational matrices. Rónyai [Ró1] has shown that distinguishing these two cases is essentially as hard as deciding quadratic residuosity modulo composite numbers. On the other hand, in the previous chapter, an ff-algorithm was obtained to compute the dimension of the minimal one-sided ideals of simple algebras over $\mathbb{Q}$. The method is based on finding maximal orders in algebras. It is natural to ask whether maximal orders can help in finding zero divisors.

In this section we settle this problem in the affirmative for a 4-dimensional simple non-commutative algebra $\mathcal{A}$ over $\mathbb{Q}$. Assume that $\mathcal{A}$ contains zero divisors. Then every maximal order in $\mathcal{A}$ is known to be isomorphic to $\mathrm{M}_2(\mathbb{Z})$, the ring of 2×2 integer matrices (cf. [Re], Theorem 21.6). Assume that a maximal order Λ is given by the set of structure constants $(c_{ij}^k)_{i,j,k=1}^4$ with respect to a basis a_1, a_2, a_3, a_4 of its additive group: $a_i a_j = \sum_{k=1}^4 c_{ij}^k a_k$ for every $i, j \in \{1,2,3,4\}$. It is clear that $c_{ij}^k \in \mathbb{Z}$. Our aim is to find a zero divisor in Λ.

For an element $a \in \mathcal{A}$, $\mathrm{tr}(a)$, the reduced trace (cf. the previous chapter) of a can be computed by dividing the trace of the left or right action of a on $\mathcal{A}$ by 4. Recall that $\mathrm{tr}(a)$ is the trace of a as a 2×2 matrix. We know that the bilinear trace form $\langle\ ,\ \rangle : \mathcal{A}\times\mathcal{A} \to \mathbb{Q}$ defined by $\langle a, b\rangle := \mathrm{tr}(ab)$ is a symmetric non-degenerate bilinear function on $\mathcal{A}$, and takes integer values on Λ. Also, for the discriminant we have $\mathrm{disc}(\Lambda) = 1$. This can be immediately seen by taking a standard basis of Λ, but it also follows from the theory of maximal orders. In fact, for $\mathcal{A} \cong \mathrm{M}_n(\mathbb{Q})$ checking the condition $\mathrm{disc}(\Lambda) = 1$ can be used to ensure the maximality of the order Λ: smaller orders have discriminant larger than 1.

Observe that for a 2×2 matrix a, such that $\mathrm{tr}(a) = 0$, the characteristic polynomial of the matrix a is $x^2 - \frac{1}{2}\mathrm{tr}(a^2)$. We shall find a very special zero divisor: a nontrivial nilpotent element, i.e., $0 \neq a \in \mathcal{A}$ such that $a^2 = 0$. Consider the linear subspace $\mathcal{A}'$ of $\mathcal{A}$ consisting of elements with zero trace and the corresponding sublattice L' of Λ, defined by

$$L' := \{a \in \Lambda : \mathrm{tr}(a) = 0\}.$$

Finding a nilpotent element of $\mathcal{A}$ is equivalent to finding an isotropic element in $\mathcal{A}'$ with respect to the bilinear trace form. Since nilpotent elements do exist, the form $\langle\ ,\ \rangle$ must be indefinite. It is easy to check that L' is a three dimensional lattice and its discriminant is equal to 2. We can find a three element basis a_1', a_2', a_3' of L' using the algorithm in [Fr].

Now we use the reduction algorithm of the previous section to compute a reduced basis (b_1, b_2, b_3) of L' with respect to the trace form. If we search for an isotropic element in the form $a = x_1 b_1 + x_2 b_2 + x_3 b_3$ then the task is equivalent to finding a nontrivial integer solution of $\sum_{i,j=1}^{3} \mathrm{tr}(b_i b_j) x_i x_j = 0$, an equation with coefficients of small size. We can make the computation more explicit: Assume that b_1, b_2, b_3 is a reduced basis, such that $b_2(1)$ and $b_3(2)$ are anisotropic. Then inequality $(2')$ of the previous section is satisfied with $c = \frac{1}{2}$, thus

$$|\langle b_1^*, b_1^*\rangle| \le 2|\langle b_2^*, b_2^*\rangle| \le 4|\langle b_3^*, b_3^*\rangle|.$$

It follows that $|\langle b_1^*, b_1^*\rangle|^3 \le 8\mathrm{disc}(L') = 16$, and since $b_1^* = b_1$ is an integer matrix with characteristic polynomial $x^2 - \frac{1}{2}\langle b_1, b_1\rangle$, we obtain $|\langle b_1^*, b_1^*\rangle| = 2$. Now we have

$$8 = 4\mathrm{disc}(L') = |\langle b_1^*, b_1^*\rangle|^3 |\langle b_2^*, b_2^*\rangle||\langle b_3^*, b_3^*\rangle| \ge |\langle b_1^*, b_1^*\rangle|^2 |\langle b_2^*, b_2^*\rangle|^2 = \mathrm{disc}(L_2')^2,$$

where L_2' is the sublattice generated by b_1 and b_2. On the other hand $|\langle b_1^*, b_1^*\rangle| \le 2|\langle b_2^*, b_2^*\rangle|$ implies

$$\mathrm{disc}(L_2')^2 = |\langle b_1^*, b_1^*\rangle|^2 |\langle b_2^*, b_2^*\rangle|^2 \ge \frac{1}{4}|\langle b_1^*, b_1^*\rangle|^4 = 4.$$

We infer that $\mathrm{disc}(L') = 2$, which gives $|\langle b_2^*, b_2^*\rangle| = 1$ and

$$|\langle b_3^*, b_3^*\rangle| = \frac{\mathrm{disc}(L')}{|\langle b_1^*, b_1^*\rangle||\langle b_2^*, b_2^*\rangle|} = 1.$$

If $\langle b_2^*, b_2^*\rangle = -\langle b_3^*, b_3^*\rangle$, then $b_2^* + b_3^*$ is isotropic. If $\langle b_2^*, b_2^*\rangle = \langle b_3^*, b_3^*\rangle$, then $\langle b_1^*, b_1^*\rangle = -2\langle b_2^*, b_2^*\rangle$ (otherwise $\langle\ ,\ \rangle$ would be definite) and in that case $b_1^* + b_2^* + b_3^*$ is isotropic. Thus in any case, we can find a nontrivial nilpotent element $a \in \mathcal{A}$ as one of the vectors

$$b_2(1),\ b_3(2),\ b_2^* + b_3^*,\ \text{or } b_1^* + b_2^* + b_3^*.$$

Since a is a nontrivial zero divisor in $\mathcal{A}$, $\mathcal{A}a = \{ba | b \in \mathcal{A}\}$ is a two-dimensional left ideal of $\mathcal{A}$ and the lattice $\Lambda a = \{ba | b \in \Lambda\}$ has rank two. Using [Fr], a $\mathbb{Z}$-basis v_1, v_2 of Λa can be found. Now we are given an isomorphism $\Lambda \cong \mathrm{M}_2(\mathbb{Z})$ by mapping $b \in \Lambda$ to the matrix of the left action of a on Λa. We have proved

Theorem 7.2.1 *Given a maximal order in an algebra isomorphic to* $\mathrm{M}_2(\mathbb{Q})$*, i.e., a ring* Λ *isomorphic to* $\mathrm{M}_2(\mathbb{Z})$ *by its integer structure constants* $\{c_{i,j}^k :\ i, j, k = 1, 2, 3, 4\}$*, we can find a nontrivial zero divisor and give an explicit isomorphism between* Λ *and* $\mathrm{M}_2(\mathbb{Z})$ *in time polynomial in the size of the structure constants.*
□

Note that since we are looking for isotropic elements, we do not need the whole power of the reduction algorithm, we can stop as soon as an isotropic vector is found. In particular, no reduction is necessary if the basis appears to be singular.

Chapter 8

Problems for further research

We conclude with some open problems related to the results in this dissertation.

8.1 Faster radical algorithms

The performance of decomposition algorithms for semisimple algebras (and modules) can be improved (see [Eb1, EG]) with the aid of choosing random (or random-like) elements, such as *splitting elements* defined in Chapter 5. The main idea of these techniques is to perform computations in *substructures* related to the elements. This usually reduces the size (the number of variables and equations) of linear systems to solve. No such improvements have been developed yet for the radical algorithms of [Di], [Ró2], and Chapter 3. It would be interesting to analyze the possibility of new radical algorithms based on computing first the radical of certain subalgebras of restricted structure, such as a Cartan subalgebra (in the Lie algebra of our associative algebra). Note that Cartan subalgebras can be efficiently computed either with the randomized or even with the deterministic method of Chapter 5.

8.2 Decomposition of simple algebras

The complexity of finding the irreducible representations of algebras over global fields remains open. Note that since we have (partly randomized) polynomial time algorithms for isolating the radical and decomposing the semisimple part, the problem is essentially equivalent to finding minimal one-sided ideals in simple algebras. Efficient algorithms for this problem would also lead to new constructive methods in *class field theory*.

Explicit isomorphism of simple algebras Assume that our algebra $\mathcal{A}$ is in fact isomorphic to $\mathrm{M}_m(K)$. Note that the ff-algorithms of Chapter 6 can be used to verify this property of $\mathcal{A}$. It is easy to see that finding a minimal one-sided ideal of $\mathcal{A}$ is equivalent to *computing an isomorphism* $\mathcal{A} \cong \mathrm{M}_m(K)$. Furthermore, using elementary ideas from the theory of *Brauer groups* (cf. [Re], Chapters 7 and 8), it is not difficult to show that this problem is in fact equivalent to computing an isomorphism between two arbitrary central simple algebras.

The results of Chapter 6, concerning the decision version of the problem, suggest the following question: how is the complexity of the explicit isomorphism problem for simple algebras related to factoring integers and polynomials? Does there exist a polynomial time ff-algorithm to solve the former problem over algebraic number fields? Does there exist a randomized polynomial time method over global function fields?

Quaternion algebras Recall from Chapter 7 that we have an affirmative answer in the special case of a *quaternion algebra* (a four dimensional central simple algebra) $\mathcal{A}$ over $\mathbb{Q}$: finding a zero divisor in $\mathcal{A}$ is (up to a polynomial time reduction) equivalent to finding a maximal order in $\mathcal{A}$, which is essentially as hard as factoring integers. It would be already interesting to generalize this result to *quaternion algebras over arbitrary global fields.* Note that it seems to be possible to generalize the method of Chapter 7 to the case where the ground field is $\mathbb{F}_q(X)$.

Rational group representations and group rings Recall from Chapter 6 that we have a polynomial time algorithm for computing the Schur indices of a finite group G given by multiplication table. In essence, this is possible since the prime factors of the discriminant of the integral group ring $\mathbb{Z}G$ are small. It does not seem to be hopeless to use the very special structure of the order $\mathbb{Z}G$ in efficient algorithms for constructing the irreducible matrix representations of G over $\mathbb{Q}$. Note that the structure of integral group rings is under intensive investigation. Let us mention here the famous conjecture of *Zassenhaus*: if G and H are finite groups such that the group rings $\mathbb{Z}G$ and $\mathbb{Z}H$ are isomorphic, then G and H are isomorphic groups. A related computational problem is testing isomorphism of integral group rings (or some more general orders).

Explicit isomorphism over the algebraic closure There are efficient algorithm for constructing *absolutely irreducible* representations of associative algebras over *extensions* of the ground field, cf. [BR, Eb1, Eb3]. Assume that $\mathcal{A}$ is a central simple algebra of dimension m^2 over the number field K. For extension L of K, the algebra $\mathcal{A}_L = \mathcal{A} \otimes_K L$ is a central simple L-algebra. Note that $\mathcal{A}_L$ has the same multiplication table as $\mathcal{A}$, but the structure constants are considered as elements of L. This construction makes possible a kind of symbolic computation over the field $\overline{\mathbb{Q}}$ of *algebraic numbers*. In this model even if the algebra $\mathcal{A}$ is inputted over a small field, we are interested in the structure of the algebra $\mathcal{A}_{\overline{\mathbb{Q}}}$. The results of [BR, Eb1, Eb3] can be formulated as follows. An extension L of K together with an isomorphism $\mathcal{A}_L \cong \mathrm{M}_m(L)$ can be computed in polynomial time.

It would be interesting to obtain analogous results for simple *Lie algebras*. In the paper [Gra] a polynomial time method is proposed to decompose a semisimple Lie algebra over an algebraic number field into a direct sum of simple ideals and to determine the type (the isomorphism class over $\overline{\mathbb{Q}}$) of each simple component. Recall that there are four infinite series ($A_l,\ B_l,\ C_l,\ D_l$) of isomorphism classes of simple Lie algebras over $\overline{\mathbb{Q}}$. In addition, in low dimensions there exist five exceptional ($E_6,\ E_7,\ E_8,\ F_4,\ G_2$) classes. It would be important to develop an efficient algorithm for *constructing an explicit isomorphism* of a simple Lie algebra with one of the standard algebras over a "small" extension of the ground field.

Bibliography

[Bab] L. Babai, Deciding finiteness of matrix groups in Las Vegas polynomial time, *Proc. 3rd ACM-SIAM Symp. on Discrete Algorithms, (1992), 33-40.*

[BBR] L. Babai, R. M. Beals, D. Rockmore, Deciding finiteness of matrix groups in deterministic polynomial time, *Israel J. Math., to appear.*

[BR] L. Babai, L. Rónyai, Computing irreducible representations of finite groups, *Mathematics of Computation 55, 192 (1990), 705-722.*

[BBCIL] L. Babai, R. Beals, J.-Y. Cai, G. Ivanyos, E. M. Luks, Multiplicative equations over commuting matrices, *Proc. 7th ACM-SIAM Symp. on Discrete Algorithms, (1996), 498-507.*

[Bach] E. Bach, Fast algorithms under the extended Riemann Hypothesis: a concrete estimate, *Proc. 14th ACM STOC, (1982), 290–295.*

[Bar] E. H. Bareiss, Sylvester's identity and multistep integer-preserving Gaussian elimination, *Mathematics of Computation 103, (1968), 565–578.*

[Bas] J. R. Bastida, Field Extensions and Galois Theory, in *Rota, G-C. (ed), Encyclopedia of Mathematics and Its Applications, Vol. 22. Cambridge University Press and Addison-Wesley, 1984.*

[Bea] R. Beals, Algorithms for matrix groups and the Tits alternative, *Proc. 36th IEEE FOCS, (1995), 593-602.*

[Ber1] E. R. Berlekamp, Factoring polynomials over finite fields, *Bell System Technical Journal 46, (1967), 1853-1859.*

[Ber2] E. R. Berlekamp, Algebraic Coding Theory. *McGraw-Hill, 1968.*

[Ber3] E. R. Berlekamp, Factoring polynomials over large finite fields, *Mathematics of Computation 24, (1970), 713-715.*

[BKS] R. E. Beck, B. Kolman, I. N. Stewart, Computing the structure of a Lie algebra, in *Computers in Non-associative Rings and Algebras, Academic Press, 1977.*

[Ca] J. W. S. Cassels, Rational Quadratic Forms, in *L.M.S. Monographs, Academic Press, 1978.*

[Ch] A. L. Chistov, Algorithm of polynomial complexity for factoring polynomials and finding the components of varieties in subexponential time, *Zapiski Nauchnykh Seminarov LOMI 137, (1984), 124–188.* English translation: *J. Soviet Math. 34, (1986), 1838–1882.*

[CG] A. L. Chistov, D. Yu. Grigoryev, Polynomial time factoring of the multivariable polynomials over a global field, *LOMI preprint E-5-82, Leningrad, 1982.*

[CIW] A. M. Cohen, G. Ivanyos, D. B. Wales, Finding the radical of an algebra of linear transformations, *4th Int. Symp. on Effective Methods in Algebraic Geometry, (1996).* To appear in J. Pure and Applied Algebra.

[CLZ] J. Cai, R. J. Lipton, Y. Zalcstein: The complexity of the membership problem for 2-generated commutative semigroups of rational matrices, *Proc. 35th IEEE FOCS, (1994), 135–142.*

[Di] L. E. Dickson, Algebras and Their Arithmetics, *University of Chicago, 1923.*

[Eb1] W. M. Eberly, Computations for Algebras and Group Representations, *Ph. D. thesis, Dept. of Computer Science, University of Toronto, 1989.*

[Eb2] W. M. Eberly, Decomposition of algebras over finite fields and number fields, *Computational Complexity 1, (1991), 179-206.*

[Eb3] W. M. Eberly, Decompositions of algebras over **R** and **C**, *Computational Complexity 1, (1991), 207-230.*

[EG] W. M. Eberly, M. Giesbrecht, Efficient decomposition of associative algebras, *Proc. ISSAC'96, (1996), 170–178.*

[Ed] J. Edmonds, System of distinct representatives and linear algebra, *Journal of Research of the National Bureau of Standards 718, 4, (1967), 241-245.*

[FR] K. Friedl, L. Rónyai, Polynomial time solution of some problems in computational algebra, *Proc. 17th ACM STOC, (1985), 153-162.*

[Fr] M. A. Frumkin, Polynomial time algorithms in the theory of linear diophantine equations, *Proc. FCT'76, (1976), 386-392.*

[FSh] A. Fröhlich, J. C. Shepardson, Effective procedures in field theory, *Royal Soc. London, Phil. Trans. A 248, (1955-56), 407-432.*

[Ge1] G. Ge, Algorithms Related to Multiplicative Representations of Algebraic Numbers, *Ph. D. thesis, Math Dept, U. C. Berkeley, 1993.*

[Ge2] G. Ge, Testing equalities of multiplicative representations in polynomial time, *Proc. 34th IEEE FOCS, (1993), 422–426.*

[GIKR] W. A. de Graaf, G. Ivanyos, A. Küronya, L. Rónyai, Computing Levi decompositions, *Applicable Algebra in Engineering, Communication and Computing, to appear.*

[GIR] W. A. de Graaf, G. Ivanyos, L. Rónyai, Computing Cartan subalgebras of Lie algebras, *Applicable Algebra in Engineering, Communication and Computing 7, (1996), 71-90.*

[GM] S. Goldwasser, S. Micali, Probabilistic encryption and how to play mental poker keeping secret all partial information, *Proc. 14th ACM STOC, (1982), 365-378.*

[GMT] P. Gianni, V. Miller, B. Trager, Decomposition of algebras, *Lecture Notes in Computer Science 358, eds. G. Goos and J. Hartmanis, Springer-Verlag 1988, 300-308.*

[Gra] W. A. de Graaf, Calculating the structure of a semisimple Lie algebra, *4th Int. Symp. on Effective Methods in Algebraic Geometry, (1996).* To appear in J. Pure and Applied Algebra.

[Gri] D. Grigoriev, Factorization of polynomials over finite field and the solution of systems of algebraic equations, *Zapiski Nauchnykh Seminarov LOMI 137, (1984), 20–79.* English translation: *J. Soviet Math. 34, (1986), 1762–1803.*

[Ha] M. Harada, Hereditary orders, *Trans. Amer. Math. Soc. 107, (1963), 273-290.*

[HU] J. E. Hopcroft, J. D. Ullman, *Introduction to automata theory, languages and computation, Addison-Wesley, 1979.*

[Hum] J. E. Humphreys, Introduction to Lie Algebras and Representation Theory, *Springer-Verlag, New York, Heidelberg, Berlin, 1972.*

[IR] G. Ivanyos, L. Rónyai, Finding maximal orders in semisimple algebras over $\mathbf{Q}$, *Computational Complexity 3, (1993), 245-261.*

[IRSz] G. Ivanyos, L. Rónyai, Á. Szántó, Decomposition of algebras over $F_q(X_1, \ldots, X_m)$, *Applicable Algebra in Engineering, Communication and Computing 5, (1994), 71-90.*

[ISz] G. Ivanyos, Á. Szántó, Lattice basis reduction for indefinite forms and an application, *Discrete Mathemathics 153, (1996), 177-188.*

[Ja] H. Jacobinski, Two remarks about hereditary orders, *Proc. Amer. Math. Soc. 28, (1971), 1-8.*

[Jac] N. Jacobson, Lie Algebras, *Dover, New York, 1979.*

[KB] R. Kannan, A. Bachem, Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix, *SIAM J. on Computing 4, (1979), 499–507.*

[Ke] A. Kertész, Lectures on Artinian rings, *Akadémiai Kiadó, Budapest 1987.*

[Kn] D. E. Knuth, The art of computer programming, Vol. 2: Seminumerical algorithms, *Addison-Wesley 1981.*

[La] S. Landau, Factoring polynomials over algebraic number fields, *SIAM J. on Computing 14, (1985), 184-195.*

[Len] A. K. Lenstra, Factoring polynomials over algebraic number fields, *Proc. EUROCAL, Lect. Notes of Comp. Sci. 162, Springer-Verlag 1983, 245–254.*

[LLL] A. K. Lenstra, H. W. Lenstra Jr., L. Lovász, Factoring polynomials with rational coefficients, *Math. Annalen 261, (1982), 515-534.*

[Lo] L. Lovász, An Algorithmic Theory of Numbers, Graphs and Convexity, *SIAM, 1986.*

[Mi] K. A. Mihaĭlova, The occurrence problem for a direct product of groups, *Dokl. Akad. Nauk 119, (1958), 1103–1105.*

[Pie] R. S. Pierce, Associative Algebras, *Springer-Verlag, 1982.*

[Po1] M. Pohst, A modification of the LLL reduction algorithm, *J. Symbolic Computation 4, (1987), 123-127.*

[Po] M. E. Pohst, Three principal tasks of computational algebraic number theory, *Number Theory and Applications, Proc. NATO Advanced Study Inst., Kluwer Academic Publishers, 1989, 123-133.*

[Ra] M. O. Rabin, Digitalized signatures and public-key functions as intractable as factorization, *MIT/LCS TR-212, Technical Memo MIT, 1979.*

[RWZ] D. Rand, P. Winternitz, H. Zassenhaus, On the Identification of a Lie algebra given by its structure constants. I. Direct decompositions, Levi decompositions, and nilradicals, *Linear Algebra and its Applications 109, (1988), 197–246.*

[Re] I. Reiner, Maximal Orders, *Academic Press, 1975.*

[Ró1] L. Rónyai, Zero divisors in quaternion algebras, *Journal of Algorithms 9, (1988), 494-506.*

[Ró2] L. Rónyai, Computing the structure of finite algebras, *J. Symbolic Computation 9, (1990), 355-373.*

[Ró3] L. Rónyai, Algorithmic properties of maximal orders in simple algebras over **Q**, *Computational Complexity 2, (1992), 225-243.*

[Ró4] L. Rónyai, Computations in associative algebras, *Groups and Computation, DIMACS Series, 11, American Mathematical Society, 1993, 221–243.*

[Ró5] L. Rónyai, A deterministic method for computing splitting elements in semisimple algebras over Q, *Journal of Algorithms 16, (1994), 24-32.*

[Sch] J. T. Schwartz, Fast probabilistic algorithms for verification of polynomial identities, *Journal of ACM 27, (1980), 701-717.*

[Z] R. E. Zippel, Propabilistic algorithms for sparse polynomials, *Proc. EUROSAM '79, Lect. Notes in Comp. Sci. 72, Springer, 1979, 216-226.*

[Win] D. J. Winter, Abstract Lie Algebras, *The MIT Press, 1972.*

Appendix: Proof of Theorem 7.1.1

I. First we show that the algorithm requires a polynomial number of arithmetical operations. This is clear for finding the initial nonsingular basis (denoted by $(c_1, \ldots, c_n)$). We have already seen that the quantity

$$D(b_1, \ldots, b_n) := \prod_{k=1}^{n} \prod_{j=1}^{k} |\langle b_j^*, b_j^* \rangle| = \prod_{k=1}^{n} \mathrm{disc}(L_k) \in \mathbb{N}.$$

remains unchanged during making the basis proper and reduces by a factor less than $\frac{3}{4}$ in every round, where a change is due to failure at tests (2a) or (2b) of reducedness. After p such rounds, we have

$$1 \leq D(b_1, \ldots, b_n) \leq (\frac{3}{4})^p D(c_1, \ldots, c_n),$$

and hence

$$p \leq \frac{1}{\log_2 \frac{4}{3}} \log_2 D(c_1, \ldots, c_n) \leq 3 \log_2 D(c_1, \ldots, c_n).$$

We conclude that the number of arithmetical operations is really bounded by a polynomial of the input size.

II. It remains to prove that the size of the numbers we work with during the run of the algorithm is also bounded by a polynomial of the input size (the dimension n, and the sizes of $\langle a_i, a_j \rangle$).

Lemmas A.1 and A.2 below are related to the procedure making a basis proper. They establish bounds on the coordinates of vectors in terms of the values of $\langle\ ,\ \rangle$ on the starting basis. Lemma A.3 provides a universal bound on the values of $\langle\ ,\ \rangle$ on an arbitrary proper basis produced by the algorithm.

Lemma A.1 *For any nonsingular basis $(b_1, \ldots, b_n)$ of L, writing $b_i = b_i^* + \sum_{j=1}^{i-1} \mu_{ij} b_j^*$ for $i = 1, \ldots, n$,*

$$|\mu_{ij}| \leq (nA)^n \quad \textit{and either } \mu_{ij} = 0 \textit{ or } \quad |\mu_{ij}| \geq (nA)^{-n}$$

where $A := \max_{k,l} |\langle b_k, b_l \rangle|$.

Proof In the arguments below we use Hadamard's bound on the determinant.

For $i > 1$ we have

$$\begin{aligned} |\langle b_i^*, b_i^* \rangle| &= \frac{|\prod_{j \leq i} \langle b_j^*, b_j^* \rangle|}{|\prod_{j < i} \langle b_j^*, b_j^* \rangle|} = \frac{|\det\left(\langle b_k, b_l \rangle\right)_{k,l \leq i}|}{|\det\left(\langle b_k, b_l \rangle\right)_{k,l < i}|} \\ &\geq |\det\left(\langle b_k, b_l \rangle\right)_{k,l<i}|^{-1} \geq ((i-1)A)^{-(i-1)} \geq (nA)^{1-n}, \end{aligned}$$

implying for all $1 \leq i < j$ that

$$|\mu_{ij}| = \frac{|\langle b_i, b_j \rangle|}{|\langle b_j^*, b_j^* \rangle|} \leq A(nA)^{n-1}.$$

Similarly,

$$|\langle b_i^*, b_i^* \rangle| \leq |\det\left(\langle b_k, b_l \rangle\right)_{k,l \leq i}| \leq (iA)^i \leq (nA)^n,$$

thus if $\langle b_i, b_j \rangle \neq 0$ then $|\mu_{ij}| \geq (nA)^{-n}$. □

Lemma A.2 *Suppose that the coordinates of the vectors b_i of a nonsingular basis $(b_1, \ldots, b_n)$ of L with respect to a fixed other basis $(d_1, \ldots, d_n)$ are bounded by C in absolute value. Then the coordinates with respect to $(d_1, \ldots, d_n)$ of the vectors occurring in the course of the procedure making the basis $(b_1, \ldots, b_n)$ proper (in particular those of the resulting proper basis) will be at most*

$$2^{n^2} B^n C,$$

where $B := (nA)^n$ with $A = \max_{kl} |\langle b_k, b_l \rangle|$.

Proof First we remark that the change of b_j to $b_j - mb_k$ affects neither the orthogonalization of the basis nor the coefficients μ_{st} for values of s other than j. We also observe that when the procedure pivots the element b_n, $(b_1, \ldots, b_{n-1})$ is already a proper basis of L_{n-1}, and by induction we can assume that for $k < n$ the absolute values of the coordinates of b_k are at most $2^{(n-1)^2} B^{n-1} C$. We prove by induction that $B2^t$ is an integer upper bound for $|\mu_{nk}|$ $(k = 1, \ldots, n-1)$ after the t'th change of b_n. For $t = 0$ this follows from the first inequality of Lemma A.1. Assume that after the $t-1$'st step $B2^{t-1}$ is an upper bound for $|\mu_{nk}|$ $(k = 1, \ldots, n-1)$. In the t'th step we perform $b_n^{(t)} := b_n^{(t-1)} - m^{(t)} b_{n-t} = \sum_{k=1}^{n} \mu_{nk}^{(t)} b_k^*$, where $|m^{(t)} - \mu_{n,n-t}^{(t-1)}| \leq \frac{1}{2}$. Observe that we just achieve $|\mu_{n,n-t}^{(t)}| \leq \frac{1}{2}$, while for $k > n-t$ we already have $|\mu_{nk}^{(t)}| = |\mu_{nk}^{(t-1)}| \leq \frac{1}{2}$. On the other hand, from the induction hypothesis and the integrality of $B2^{t-1}$ we obtain $|m^{(t)}| \leq B2^{t-1}$ and hence $|\mu_{nk}^{(t)}| = |\mu_{nk}^{(t-1)} - m^{(t)} \mu_{n-t,k}| \leq B2^{t-1} + \frac{1}{2} B2^{t-1} < B2^t$ $(k = 1, \ldots, n-t-1)$. We used here the fact that $k < n$, $n-t < n$ and therefore we already have $|\mu_{n-t,k}| < \frac{1}{2}$. After the t'th step $(1 \leq t \leq n-1)$, b_n changes to

$$b_n^{(t)} = b_n - (m^{(1)} b_{n-1} + \ldots + m^{(t)} b_{n-t}).$$

Using the bound $|m^{(s)}| \leq B2^{s-1}$ obtained above and the induction hypothesis on $b_1, \ldots, b_{n-1}$, we infer the following upper bound on absolute values of the coordinates of $b_n^{(t)}$:

$$2^{(n-1)^2} B^{n-1} C (1 + B + \cdots + B2^t) \leq 2^{(n-1)^2} B^{n-1} C B 2^{t+1} < 2^{n^2} B^n C,$$

as claimed. □

Lemma A.3 *For any proper basis $(b_1, \ldots, b_n)$ produced by the algorithm, we have*

$$\max_{i,j} |\langle b_i, b_j \rangle| \leq nD(c_1, \ldots, c_n)$$

where $(c_1, \ldots, c_n)$ is the initial nonsingular basis.

Proof If we write $b_i = \sum_{j=1}^{i} \mu_{i,j} b_j^*$ $\quad (i = 1, \ldots, n)$ then $|\mu_{i,j}| \leq 1$ $(i = 1 \ldots n,\ j = i \ldots n)$, thus we have

$$|\langle b_i, b_k \rangle| = \left| \sum_{j=1}^{\min(i,k)} \mu_{ij} \mu_{kj} \langle b_j^*, b_j^* \rangle \right| \leq \sum_{j=1}^{\min(i,k)} |\langle b_j^*, b_j^* \rangle|.$$

We have seen in the first part of the proof that the quantity $D(b_1, \ldots, b_n)$ is decreasing during the algorithm, thus

$$|\langle b_j^*, b_j^* \rangle| = \frac{\mathrm{disc}(L_j)}{\mathrm{disc}(L_{j-1})} \leq \mathrm{disc}(L_j) \leq D(c_1, \ldots, c_n),$$

and the claim follows easily. □

Now we can estimate the size of the numbers we work with. The maximum of $\{|\langle a_i, a_j\rangle| : i, j = 1, \dots, n\}$ is denoted by A_0. Let $(c_1, \dots, c_n)$ be the starting nonsingular basis. We know that either $c_i = a_i$ or $c_i = a_i + \epsilon a_j$ for some $j > i$ and $\epsilon = \pm 1$. From this we infer that $|\langle c_i, c_j\rangle| \leq 4A_0$ and therefore (using Hadamard's bound on the n discriminants)

$$D(c_1, \dots, c_n) \leq (4nA_0)^{n^2}.$$

We denote by C_1 the maximum of the absolute values of the coordinates of the first proper basis $(b_1, \dots, b_n)$. By Lemma A.2 with $B_0 := (4nA_0)^n$ we have

$$C_1 \leq 2 \cdot 2^{n^2} B_0^n.$$

This bound is also valid for the intermediate vectors of the procedure making the initial nonsingular basis $(c_1, \dots, c_n)$ proper.

Lemma A.3 gives an absolute bound for values of $\langle\ ,\ \rangle$ on the proper bases produced by the algorithm:

$$|\langle b_i, b_j\rangle| \leq nB_0^n =: E.$$

A swap performed due to failure at test (2a) does not affect the set of the coordinates. In case of failure at (2b), we perform $b_i := b_i + zb_{i+1}$ for some i, where $|z| \leq \frac{1}{2}(nE)^n + 1 < (nE)^n$. Since z is an integer, $|z+1| \leq (nE)^n$. For the new values we have $|\langle b_i, b_j\rangle| \leq (1 + |z|)^2 E$, thus $A := E(nE)^{2n} < (nE)^{2n+1}$ is an upper bound for the absolute values of $\langle\ ,\ \rangle$ before each entry to the procedure maintaining properness. Now let C_k stand for the maximum of the absolute values of the coordinates of the elements of the k'th proper basis. With $B := (nA)^n$ we have

$$C_k \leq C_{k-1}(nE)^n B^n 2^{n^2}.$$

The algorithm terminates in at most $p = 3\log_2 D(c_1, \dots, c_n)$ iteration steps, and

$$C_p \leq 2(nEB)^{np} 2^{n^2 p}.$$

C_p is an upper bound for all the coordinates computed in our algorithm. For the size of the coordinates this yields the bound

$$\log_2 C_p = 1 + n^2 p + np\log_2(2nE(1+B)) = O((n^7)\log^2(n)\ \log^2(A_0)),$$

since $p = O(n^2(\log(n) + \log(A_0)))$, $\log(E) = O(n^2(\log(n) + \log(A_0)))$ and $\log(B) = O(n^4(\log(n) + \log(A_0)))$.

□