

Polynomial time algorithms for modules over finite dimensional algebras

Alexander Chistov ^{*} Gábor Ivanyos [†] Marek Karpinski[‡]

May 13, 1997

Abstract

We present polynomial time algorithms for some fundamental tasks from representation theory of finite dimensional algebras. These involve testing (and constructing) isomorphisms of modules as well as expressing of modules as direct sums of indecomposable modules. Over number fields the latter task seems to be difficult, therefore we restrict our attention to decomposition over finite fields and over the algebraic or real closure of number fields.

The module isomorphism problem can be reformulated as follows. Let A_1, \dots, A_m and A'_1, \dots, A'_m be two families of $n \times n$ -matrices with entries from the field K . The task is to find a nonsingular $n \times n$ -matrix X with entries from K such that $XA_iX^{-1} = A'_i$ for all $1 \leq i \leq m$ (if such a matrix exists). In the case when K is the field of the real algebraic numbers, we propose a method for the variant where the matrix X is required to be orthogonal.

^{*}St. Petersburg Institute for Informatics and Automation of the Academy of Sciences of Russia, 14th line 39, St. Petersburg 199178, Russia. Research partially supported by the Volkswagen-Stiftung, Program on Computational Complexity.

[†]Computer and Automation Institute, Hungarian Academy of Sciences, Lágymányosi u. 11, H-1111 Budapest, Hungary. Research supported in part by OTKA Grant 016503, AKP Grant 96-350/24, and EC Grant ALTEC-KIT.

[‡]Department of Computer Science, University of Bonn, 53117 Bonn and International Computer Science Institute, Berkeley, California. Research partially supported by the DFG Grant KA 673/4-1, by the ESPRIT BR Grants 7097 and EC-US 030 and by the Max-Planck Research Prize.

1 Introduction

In this paper we discuss the computational complexity of some principal tasks from representation theory of finite dimensional algebras such as testing and constructing isomorphisms of modules and decomposing modules as direct sums of indecomposable submodules. These problems arise naturally in computations involving several matrices. Holt and Rees [8] presents a randomized algorithm (an improvement of the *meat-axe* method by Parker and Norton [10]) for finding proper submodules of modules over finite algebras. Except the case of modules having a very special structure, the algorithm finds a proper submodule with high probability. Holt and Rees also show how to apply their procedure to other related tasks, such as a special case of the module isomorphism problem. It is worth to note that the method is very efficient in practice as well. The algorithm is implemented into the computer algebra systems GAP and MAGMA. In this paper we allow more general ground fields, but we are concerned with *theoretical complexity* and give *polynomial time* algorithms. Furthermore, most of our algorithms are *deterministic*.

Module isomorphism can be translated to the following conjugacy problem. Assume that we are given two collections A_1, \dots, A_m and A'_1, \dots, A'_m of $n \times n$ matrices with entries from the field K . Our task is to find (if exists) a nonsingular matrix $X \in \text{GL}_n(K)$ such that

$$XA_iX^{-1} = A'_i$$

for every $i \in \{1, \dots, m\}$. A natural approach to this problem is to consider the set V of $n \times n$ matrices X satisfying

$$XA_i = A'_iX.$$

This condition is equivalent to a system of homogeneous linear equations in the entries of the matrix X , whence V is a linear subspace of $\text{Mat}_n(K)$. Obviously, the conjugacy problem is equivalent to finding a nonsingular matrix in the subspace V . Thus, the conjugacy problem can be considered as a special case of finding nonsingular matrices in linear subspaces of $\text{Mat}_n(K)$, which was formulated by Edmonds [6]. Note that if our ground field K is sufficiently large then this problem admits an efficient randomized solution: If there exists a nonsingular matrix in the linear subspace V of $\text{Mat}_n(K)$

then a random matrix from V (i.e., a random linear combination of a basis) is nonsingular with high probability ([15, 16]). However, no deterministic polynomial time method is known to this general problem. In this paper we present a deterministic polynomial time algorithm for our particular case, i.e., the conjugacy problem.

Module decompositions correspond to common block structures of several matrices. Let A_1, \dots, A_m be matrices from $\text{Mat}_n(K)$. These matrices act naturally on the space $V = K^n$ of column vectors of length n . Assume that $0 = V_0 < V_1 < \dots < V_r = V$ is an ascending sequence of subspaces of V invariant with respect to every $A_i \in \{A_1, \dots, A_m\}$. After switching to an appropriate basis (choosing a basis of V_1 , extending it to a basis of V_2 , and iterating this procedure) the matrices A_1, \dots, A_m will have block upper triangular form. The diagonal blocks correspond to the actions of our matrices on the factor spaces V_i/V_{i-1} . A finest common block triangular form of A_1, \dots, A_m corresponds to a composition series of V , that is a finest ascending sequence of invariant subspaces.

A special case is where V decomposes as a direct sum $V_1 \oplus \dots \oplus V_m$ of common invariant subspaces. In a basis that is a union of bases of the subspaces V_i and V our matrices have block diagonal forms. A finest common block diagonal structure correspond to decomposition of V as a direct sum of indecomposable invariant subspaces.

Obviously, the subspaces invariant w. r. t. A_1, \dots, A_m are also invariant with respect to the entire matrix algebra $\Lambda \subset \text{Mat}_n(K)$ generated by the identity matrix I and A_1, \dots, A_m . This is the linear span of I and products of the form $A_{i_1} \cdots A_{i_s}$.

For the basic notions and facts related to finite dimensional associative algebras and modules the reader is referred to the book [11]. In this paper by a K -algebra we mean a finite dimensional algebra over the field K with identity element. Modules over the algebra Λ are assumed to be finite dimensional, unital left Λ -modules. (The Λ -module V is called unital if the identity element of Λ acts identically on V .) Note that this is not an essential restriction since every module can be decomposed in an obvious way as the direct sum of a unital submodule and a zero submodule.

An algebra Λ can be described by a collection of *structure constants* with respect to a basis. Assume that a_1, \dots, a_n is a K -basis of the algebra Λ . The multiplication of Λ is determined by the table of multiplication of the basis

elements. We have

$$a_i a_j = \sum_{h=1}^n \gamma_{ij}^h a_h, \quad (i, j \in \{1, \dots, n\})$$

for some elements $\gamma_{ij}^h \in K$, called the structure constants. Modules can be represented in a similar way. Let V be a Λ -module with basis v_1, \dots, v_d . The action of Λ on V is determined by the elements $\eta_{ij}^h \in K$, given as

$$a_i v_j = \sum_{h=1}^d \eta_{ij}^h v_h, \quad (i \in \{1, \dots, n\}, j \in \{1, \dots, d\}).$$

An alternative way to define an algebra together with a module would be to give a (finite) set S of linear transformations on the vector space V . The algebra Λ is the subalgebra of $\text{End}_k(V)$ generated by the set $S \cup \{\text{Id}\}$. The action of Λ on V is defined in the natural way. This representation is slightly restrictive, since V is a faithful Λ -module. On the other hand, a basis of Λ together with multiplication tables can be computed in polynomial time (at least over most of the interesting ground fields). In this paper we assume that algebras and modules are given by structure constants over K .

We intend to use methods based on computing structural invariants of algebras [7, 13, 5] to the module problems outlined above. Unfortunately, finding a nontrivial proper left ideal of Λ is already difficult (essentially as hard as factoring integers), even when Λ is a non-commutative simple algebra over \mathbb{Q} of dimension four [12]. Even worse, it is not known in general whether minimal left ideals of polynomial size exists in non-commutative simple algebras. To get around this difficulty, we restrict our attention to decomposition over finite fields and over the real or algebraic closure of number fields.

Assume that E is an extension field of K . We can consider the E -algebra $\Lambda^E = E \otimes_K \Lambda$. This corresponds to the situation where an E -algebra Λ' is given by a basis such that the structure constants fall in the subfield K . Then the K -span of the basis is a K -algebra Λ and $\Lambda' \cong E \otimes_K \Lambda$. In this situation we say that Λ' has a K -structure, Λ or Λ' is *defined over* the field K . Let L be an intermediate field: $K \leq L \leq E$. Then the L -span Λ^L of the basis is an L -algebra isomorphic to $L \otimes_K \Lambda$. The elements of Λ' which belong to Λ^L are said to be *defined over* L . The K -subspace (subalgebra, ideal, etc.) W' of Λ' is said to be *defined over* L if $W' = K \otimes_L W$ for some

L -subspace (subalgebra, ideal, etc.) W of Λ^L . The notion of (sub)modules defined over L can be introduced in a similar way.

We are interested in the cases when E is one of the following.

- $E = \mathbb{C}_0$, the field of the algebraic numbers;
- $E = \mathbb{R}_0$, the field of the real algebraic numbers.

In both cases we assume that our input is defined over an algebraic number field K . The field K is assumed to be given by a minimal polynomial $f(x)$ of a primitive element α over the prime field. In addition, we assume that an isolating rectangle (interval) is also given that distinguishes α from other (real) roots of $f(x)$. This is necessary to fix the embedding $K \leq E$.

Since K is a perfect field, the radical of algebras and modules can be defined over K . In general, there is no small common intermediate field of definition of the simple components of semisimple algebras over E [5]. Therefore we have to allow different constituents of the output of decomposition algorithms to be defined over different intermediate fields.

Note that algebras over the ground field $\overline{\mathbb{F}}_p$, the algebraic closure of the finite field \mathbb{F}_p could be also of interest. However, results over $\overline{\mathbb{F}}_p$ are rather straightforward applications of the algorithms for algebras and modules over finite fields. We leave the details to the reader.

Elements of finite fields and number fields are represented as polynomials in the defining primitive element over the prime field. Of course, the polynomial is assumed to be reduced modulo the minimal polynomial of the primitive element. We use the customary representations of the elements of the various prime fields. Substructures such as submodules, subalgebras, etc. are assumed to be represented by bases. The size of a (possibly compound) object is the total number of the bits representing the object.

Below we state the main results of the paper. We keep ourselves to the convention that the term K -algebra is reserved for finite dimensional K -algebras with identity element. Similarly, the term module is reserved for finite dimensional unital modules.

Recall that a Λ -module V is *cyclic* if V is generated by a single element, i.e, there exists an element $v \in V$ such that $\Lambda v = V$. Note that V is a cyclic module if and only if there exists an epimorphism ${}_{\Lambda}\Lambda \rightarrow V$. (By ${}_{\Lambda}\Lambda$ we denote the regular Λ -module. This is the space Λ where the elements of Λ

act by multiplication from the left. The submodules of ${}_{\Lambda}\Lambda$ are left ideals of Λ .)

Theorem 1 *Let K be either a finite field or an algebraic number field, Λ be a K -algebra and V be a Λ -module. Given Λ and V , one can decide by a deterministic polynomial time algorithm, whether the module V is cyclic. If this is the case the algorithm returns a generator of V .*

As an application, we have the following result on a generalized conjugacy problem.

Theorem 2 *Let Λ be a K -algebra, where K is a finite field or an algebraic number field. Assume that we are given two collections a_1, \dots, a_m and a'_1, \dots, a'_m of elements from Λ . We can decide in deterministic polynomial time whether there exists an element $x \in \Lambda^*$ such that $xa_i x^{-1} = a'_i$ for every $i = 1, \dots, m$, and exhibit such an element if one exists.*

A standard argument shows that the module isomorphism problem is equivalent to the conjugacy problem in the full matrix algebra $\text{Mat}_n(K)$.

Corollary 3 *Let Λ be a K -algebra, where K is a finite field or an algebraic number field. Assume that we are given two Λ -modules V and W . One can decide in deterministic polynomial time whether V and W are isomorphic, and if it is the case then construct an isomorphism between these two modules.*

We prove Theorems 1, 2, and Corollary 3 in Section 2. We remark that the analogous theorems hold if we are interested in finding generators, conjugating elements, and a module isomorphism over the algebraic or real closure E of a number field K . The results will be defined over the ground field K . We leave the details to the reader.

In Section 3 we present an efficient algorithm for a variant of the conjugacy problem. We prove

Theorem 4 *Let K be a real algebraic number field. Assume that we are given two families a_1, \dots, a_s and a'_1, \dots, a'_s of matrices from $\text{Mat}_n(K)$. One can decide in polynomial time whether there exists an orthogonal matrix $x \in O_n(\mathbb{R}_0)$ such that $xa_i x^{-1} = a'_i$ for every $i \in \{1, \dots, s\}$, and exhibit such a matrix $x \in O_n(\mathbb{R}_0)$ if one exists. (Here, parts of the output x are allowed to be defined over different real extensions of K).*

Section 4 is devoted to decomposition of modules into direct sums of irreducibles. We prove

Theorem 5 *Let K be a finite field of characteristic p , Λ be a K -algebra, and V be a Λ -module. There is a Las Vegas polynomial time algorithm that finds indecomposable submodules V_1, \dots, V_m such that $V = V_1 \oplus \dots \oplus V_m$. The same task can be done by a deterministic method running in time $(\text{input size} + p)^{O(1)}$.*

Over the (real) algebraic number we have

Theorem 6 *Let K be an algebraic number field, and E be the algebraic or real closure of K . In the latter case we require K to be a real field. Assume that Λ is a K -algebra, and V is a Λ -module. One can construct in deterministic polynomial time indecomposable submodules V_1, \dots, V_m of V^E such that $V^E = V_1 \oplus \dots \oplus V_m$. The submodules V_i are defined over the (real) algebraic number fields L_i which are extensions of K .*

2 Isomorphism of modules and the conjugacy problem

We prove Theorem 1 first for semisimple modules (Sec. 2.1). Then, in Section 2.2, we "lift" the result from the factor by the radical. Finally we show how Theorem 1 applies to the generalized conjugacy problem (Sec. 2.3) and to the module isomorphism problem (Sec. 2.4).

2.1 Finding free submodules over semisimple algebras

In this subsection Λ is a semisimple algebra over the field K and V is a Λ -module.

For every $v \in V$ we consider the module homomorphism $\phi_v : {}_{\Lambda}\Lambda \rightarrow V$ given as $\phi(x) = xv$. We define the rank $\text{rk } v$ of $v \in V$ as the rank of the linear transformation $\phi_v : \Lambda \rightarrow V$.

Recall that the annihilator $\text{Ann}_{\Lambda}(v)$ of an element $v \in V$ in Λ is a left ideal of Λ given as $\{x \in \Lambda \mid xv = 0\}$. This is the kernel of the Λ -module homomorphism $\phi_v : x \mapsto xv$, whence $\Lambda v \cong {}_{\Lambda}\Lambda / \text{Ann}_{\Lambda}(v)$. We have $\text{rk } v = \dim \text{im } \phi_v = \dim \Lambda v = \dim \Lambda - \dim \text{Ann}_{\Lambda}(v)$.

An element $v \in V$ is of maximal rank if $\text{rk } w \leq \text{rk } v$ for every $w \in V$. The following lemma suggests a method for testing whether v is of maximal rank.

Lemma 7 *Let V be a module over the semisimple K -algebra Λ . The element $v \in V$ is of maximal rank if and only if $\text{Ann}_\Lambda(v)V \subset \Lambda v$.*

Proof: Let V_1, \dots, V_s be representatives of the isomorphism classes of the simple Λ -modules. Assume that ${}_\Lambda\Lambda \cong V_1^{\mu_1} \oplus \dots \oplus V_s^{\mu_s}$ and $V \cong V_1^{\nu_1} \oplus \dots \oplus V_s^{\nu_s}$. It is easy to see that the Λ -module U is cyclic if and only if $U \cong V_1^{\kappa_1} \oplus \dots \oplus V_s^{\kappa_s}$, where $\kappa_1 \leq \mu_1, \dots, \kappa_s \leq \mu_s$. It follows that v is of maximal rank in V if and only if the submodule Λv is isomorphic to $V_1^{\min\{\mu_1, \nu_1\}} \oplus \dots \oplus V_s^{\min\{\mu_s, \nu_s\}}$.

Assume that v is not of maximal rank. Then there exists a simple Λ -module, say V_1 , such that the multiplicity κ_1 of V_1 in Λv is less than $\min\{\mu_1, \nu_1\}$. Assume further that $\text{Ann}_\Lambda(v)V \subset \Lambda v$, in other words, $\text{Ann}_\Lambda(v)$ annihilates the factor module $V/\Lambda v$. The multiplicity of V_1 in that factor module is $\nu_1 - \kappa_1 > 0$, therefore $\text{Ann}_\Lambda(v)$ annihilates the module V_1 as well. But $\text{Ann}_\Lambda(V_1)$ is the ideal of Λ complementary to the ideal generated by the minimal left ideals isomorphic to V_1 . This is a contradiction since the multiplicity of V_1 in $\text{Ann}_\Lambda(v)$ is $\mu_1 - \kappa_1 > 0$. The "if" part is proved.

We give a proof of the "only if" part that will be useful in algorithms. Since Λ is semisimple there exist a left ideal L in Λ complementary to $\text{Ann}_\Lambda(v)$: ${}_\Lambda\Lambda = L \oplus \text{Ann}_\Lambda(v)$. Similarly, there exists a submodule V' of V complementary to Λv . The map ϕ_v induces an isomorphism $L \cong \Lambda v$. Assume that we have bases of Λ and V , respectively, that reflect the decompositions described above. By this we mean that the basis of Λ is a union of bases of L and $\text{Ann}_\Lambda(v)$, while the basis of V is a union of bases of Λv and V' . For every $w \in V$ we consider the block structure of the matrix of ϕ_w . We see that the matrix of ϕ_v is a regular matrix on the block corresponding to $L \times \Lambda v$, and zero outside that block. Assume that there exists element $w \in V$ such that $\text{Ann}_\Lambda(v)w$ is not a subset of Λv . Decompose w as $w = cv + w'$, where $c \in \Lambda$ and $w' \in V'$. Since $\text{Ann}_\Lambda(v)cv \subset \Lambda v$, $\text{Ann}_\Lambda(v)w' \not\subset \Lambda v$. Observe that both blocks of the matrix of $\phi_{v'}$ corresponding to Λv are zeros. It follows that the matrix of $\phi_{v+w'}$ is a block triangular matrix (both ϕ_v and $\phi_{w'}$ are zeros in the block corresponding to $\text{Ann}_\Lambda(v) \times \Lambda v$), whence the sum of the ranks of the diagonal blocks is a lower bound for $\text{rk}(v + w')$. In particular, $\text{rk}(v + w') > \text{rk}(v)$. We have proved the lemma. \square

The argument above also suggests a test of rank maximality as well as a method for incrementing the rank if it is possible. Indeed, let $v \in V$ and let v_1, \dots, v_r be a basis of V . Obviously, $\text{Ann}_\Lambda(v)V \leq \Lambda v$ if and only if $\text{Ann}_\Lambda(v)w \subset \Lambda v$ for every $w \in \{v_1, \dots, v_r\}$. We can compute the annihilator $\text{Ann}_\Lambda(v)$ and test whether $\text{Ann}_\Lambda(v)w \in \Lambda v$ for every $w \in \{v_1, \dots, v_r\}$ via solving systems of linear equations. This procedure terminates either with the conclusion that v is of maximal rank or with the first element $w \in \{v_1, \dots, v_r\}$ such that $\text{Ann}_\Lambda(v)w \not\subset \Lambda v$. We can compute a projection $\pi \in \text{End}_\Lambda(V)$ such that $\text{im } \pi = \Lambda v$ and $\pi(v) = v$ via solving a system of linear equations. We take $w' = w - \pi(w)$. The argument of the proof of the lemma shows that $\text{rk}(v + w') > \text{rk}(v)$.

This method could serve as a basic step of iteration in a procedure for finding an element $v \in V$ of maximal rank. In fact, the procedure performs polynomially many field operations. Unfortunately, over infinite ground fields, we solve systems of linear equations that depend on the previous intermediate vector v , therefore we do not have any good control over the sizes of the vectors that occur during the iteration. Over sufficiently large fields we have the following generalization of [1], Lemma 5.2.

Lemma 8 *Let V be an r -dimensional module over the semisimple K -algebra Λ and v_1, \dots, v_r be a K -basis of V . Assume that $v \in V$ is an element of non-maximal rank. Let Ω be a subset of K^* consisting of at least $\text{rk } v + 1$ elements. Then there exists a scalar $\omega \in \Omega$ and a basis element $u \in \{v_1, \dots, v_r\}$ such that $\text{rk}(v + \omega u) > \text{rk } \Lambda v$, i.e., $\dim_K \Lambda(v + \omega u) > \dim_K \Lambda v$.*

Proof: We use an argument similar to the proof of Lemma 5.2. in [1]. Let $w \in \{v_1, \dots, v_r\}$ such that $\text{Ann}_\Lambda(v)w \not\subset \Lambda v$. As in the proof of the preceding lemma, we consider decompositions ${}_\Lambda \Lambda = L \oplus \text{Ann}_\Lambda(v)$ and $V = \Lambda v \oplus V'$ as well as the related block structure of matrices of ϕ_v and ϕ_w . Let $l = \text{rk}(v)$. By choosing bases appropriately, we can achieve the situation where the matrix of ϕ_v is zero except the $l \times l$ principal minor, and the entry in position $(l + 1, l + 1)$ of the matrix of ϕ_w is nonzero. We also know that the $l \times l$ principal minor of ϕ_v has rank l . Let x be an indeterminate and $d(x)$ be the determinant of the $(l + 1) \times (l + 1)$ minor of the matrix of $\phi_{v+xw} = \phi_v + x\phi_w$. Obviously, $d(x) \in K[x]$ is of degree at most $l + 1$. Expanding the determinant at the last row one sees that the coefficient of the linear term in $d(x)$ is the determinant

of the $l \times l$ principal minor of ϕ_v . In particular, $d(x)$ is a nonzero polynomial of degree at most $l + 1$. Since $\Omega \cup \{0\} > l + 1$, there exists $\omega \in \Omega$ such that $d(\omega) \neq 0$. This implies that for such a scalar ω $\text{rk}(v + \omega w) \leq l + 1$. \square

This lemma suggests another iterative method for finding an element $v \in V$ of maximal rank, provided that our ground field K is sufficiently large. Let v_1, \dots, v_r be a basis of V and Ω be a subset of K^* of cardinality r . Initially we take $v = 0$. In each round, we compute the ranks $\text{rk}(v + \omega w)$, ($w \in \{v_1, \dots, v_r\}, \omega \in \Omega$). We replace v with the first element $v + \omega w$ such that $\text{rk}(v + \omega w) > \text{rk}(v)$. We stop if no such element exists. The procedure terminates in at most r iterations and the intermediate element v after t rounds is in the form $\omega_1 w_1 + \dots + \omega_t w_t$, where $\omega_i \in \Omega$ and $w_i \in \{v_1, \dots, v_r\}$. If K is an algebraic number field, we take $\Omega = \{1, \dots, r\}$. This gives a polynomial bound on the size of the vectors we compute with. We have proved the following.

Theorem 9 *Let V be a module over the semisimple K -algebra Λ , where K is a finite field or an algebraic number field. There is a deterministic polynomial time algorithm that finds an element $v \in V$ of maximal rank. \square*

We also have a straightforward generalization of the procedure *findfree* of the paper [1].

Theorem 10 *Let V be a module over the semisimple K -algebra Λ , where K is a finite field or an algebraic number field. There is a deterministic polynomial time algorithm that finds (free generators of) a maximal free submodule of V . \square*

2.2 Finding a single generator

In this subsection we return to the general case where V is a module over the (not necessarily semisimple) K -algebra Λ and prove Theorem 1.

We compute the radical $\text{Rad}\Lambda$ by the methods of [7] (number field case) or [13] (finite field case). Using $\text{Rad}\Lambda$, we can compute $\text{Rad}V = (\text{Rad}\Lambda)V$. We consider the action of the factor-algebra $\overline{\Lambda} = \Lambda/\text{Rad}\Lambda$ on $\overline{V} = V/\text{Rad}V$. Let $v \in V$ be an arbitrary vector and $\overline{v} = v + \text{Rad}V$. It is obvious that $\Lambda v = V$ implies $\overline{\Lambda}\overline{v} = \overline{V}$. We claim that converse also holds. The proof relies on the well known fact that elements of $\text{Rad}V$ can be omitted from

any system of Λ -module generators. Assume that \bar{v} is a generator of the $\bar{\Lambda}$ -module \bar{V} . This means that $\Lambda v + \text{Rad}V = V$. Assume that Λv is a proper submodule of V . Let M be a maximal proper submodule of V containing Λv . Since $\text{Rad}V \leq M$, we have $\Lambda V + \text{Rad}V \leq M < V$, a contradiction. We have proved the claim.

\bar{V} is a unital module over the semisimple algebra $\bar{\Lambda}$. Using the method of Theorem 9 we compute an element $\bar{v} \in \bar{V}$ of maximal rank. If $\text{rk } \bar{v} < \dim_K \bar{V}$, i.e., \bar{v} is not a generator then neither V nor \bar{V} is cyclic. On the other hand, if \bar{v} is a generator of \bar{V} then we can return any element $v \in \bar{v}$ as a generator of V . This finishes the proof of Theorem 1. \square

2.3 The general conjugacy problem

This subsection is devoted to the proof of Theorem 2.

We consider the linear subspace V of Λ given as

$$V = \{v \in \Lambda \mid va_i = a'_i v \ (i = 1, \dots, m)\}.$$

The task is equivalent to finding a unit in V . Let Λ' be the centralizer of the elements a_1, \dots, a_m :

$$\Lambda' = \{x \in \Lambda \mid xa'_i = a'_i x \ (i = 1, \dots, m)\}.$$

Λ' is a subalgebra of Λ containing 1_Λ and V is closed under multiplication by elements from Λ' from the left, i.e., V is a left Λ' -module. Let v be an arbitrary element from V . We use the linear map $\phi_v : \Lambda' \rightarrow V$ mapping x to xv . We claim that if $\Lambda^* \cap V \neq \emptyset$ then V is a cyclic Λ -module and every generator v of Λ is a unit in Λ . Indeed, let $y \in \Lambda^* \cap V$. the map ϕ_y is a Λ' module isomorphism between Λ' and V : the inverse of ϕ_y is the map $w \mapsto wy^{-1}$. In particular, V is cyclic. Let x be an arbitrary generator. Then xy^{-1} is a generator of $\Lambda'\Lambda'$, therefore xy^{-1} is a unit in Λ' , whence $xy^{-1} \in \Lambda^*$, and $y \in \Lambda^*$. The claim is proved.

We compute V and Λ' as the solution spaces of systems of linear equations. We attempt to find a generator of V by the method of Theorem 1. If V is not cyclic then the conjugacy problem admits no solution. If V is cyclic then the method of Theorem 1 also returns a generator x of V . Again, if x is not a unit then there exist no units in V at all. Otherwise we can return x . Theorem 2 is proved. \square

2.4 Module isomorphism

There is a rather obvious correspondence between the conjugacy problem in full matrix rings and the module isomorphism problem. Here we only show a reduction from the module isomorphism to the conjugacy problem. Let V and W be two unital Λ -modules. Let x_1, \dots, x_s be a set of algebra generators of Λ , e.g., a K -basis. For every $i = 1, \dots, s$ let $a_i \in \text{End}_K(V)$ be the action of x_i on V : $a_i : v \mapsto x_i v$. We define the linear transformations $b_i \in \text{End}_K(W)$ similarly. Obviously, if $V \cong W$ then $\dim V = \dim W$, therefore we may assume that $\dim V = \dim W$ and we can fix a K -isomorphism $\psi : V \rightarrow W$. For example, if V , resp. W are given by bases v_1, \dots, v_n and w_1, \dots, w_n then $\psi : v_i \mapsto w_i$ ($i = 1, \dots, n$) is a natural choice. For every $i = 1, \dots, s$, let $a'_i = \psi^{-1} \circ b_i \circ \psi : \text{End}_K(V)$. Using the particular ψ above, the matrix of a'_i (in terms of v_1, \dots, v_n) is same as the matrix of b_i in terms of w_1, \dots, w_n . A linear map $\phi \in \text{Hom}_K(V, W)$ is a K -module isomorphism if and only if the linear transformation $\eta = \phi \circ \psi \in \text{End}_K(V)$ is a unit in $\text{End}_K(V)$ and $\eta^{-1} a_i \eta = a'_i$ for every $i = 1, \dots, s$. In other words, η is a solution of a conjugacy problem in the algebra $\text{End}_K(V)$. This finishes the proof of Corollary 3.

3 Conjugacy over the orthogonal group

In this section we prove Theorem 4.

Let a_1, \dots, a_s and a'_1, \dots, a'_s be $n \times n$ matrices from $\text{Mat}_n(K)$. We would like to decide whether there exists an orthogonal matrix $x \in \text{O}_n(\mathbb{R}_0)$ such that $x a_i x^{-1} = a'_i$.

Let us denote the transpose of a matrix $a \in \text{Mat}_n(\mathbb{R}_0)$ by a^T .

We claim that there exists an orthogonal matrix $x \in \text{O}_n(\mathbb{R}_0)$ such that $x a_i x^{-1} = a'_i$ if and only if there is a regular matrix $y \in \text{GL}(\mathbb{R}_0)$ such that

$$(*) \quad y a_i y^{-1} = a'_i \text{ and } y a_i^T y^{-1} = a_i'^T, \quad i = 1, \dots, s.$$

It is straightforward to see that any orthogonal solution x to the original conjugacy problem satisfies $(*)$ as well. To prove the "if" part of the claim, assume that $y \in \text{GL}(\mathbb{R}_0)$ is a solution to $(*)$. We consider the matrix $z = y y^T$. This is a positive definite symmetric matrix. It follows that z is similar (over \mathbb{R}_0) to a diagonal matrix with positive entries, therefore there exists a matrix w in the subalgebra of $\text{Mat}_n(\mathbb{R}_0)$ generated by Id and z such that $z = w^2$.

Note that z is also a symmetric matrix. For every $i \in \{1, \dots, s\}$ we have $za'_i = yy^T a'_i = y(a'_i{}^T y)^T = y(ya_i{}^T)^T = ya_i y^T = a'_i yy^T = a_i z$, i.e., z commutes with the matrix a'_i . Since the matrix w is a polynomial of z , the same holds for the matrix w . Now we take $x = w^{-1}y$. Since w commutes with the matrices a_i , we have $xa_i x^{-1} = w^{-1}ya_i y^{-1}w = w^{-1}a'_i w = a'_i$ ($i = 1, \dots, s$). On the other hand, $xx^T = w^{-1}yy^T w^{-1} = w^{-1}w^2 w^{-1} = \text{Id}$, therefore $x \in \text{O}_n(\mathbb{R}_0)$. The claim is proved.

Since $(*)$ is a conjugacy problem with two lists, each consisting of $2m$ matrices, we can use Theorem 2 for finding a solution $y \in \text{GL}(K)$. If there is no such y , there is no one even in $\text{GL}(\mathbb{R}_0)$, therefore the conjugacy problem admits no orthogonal solution.

The only serious algorithmic problem is constructing a “square root” w of the matrix $z = yy^T$. Let $f(t) \in K[t]$ be the characteristic polynomial of z . Let ζ_1, \dots, ζ_s be the roots of $f(t)$. We know that ζ_1, \dots, ζ_s are positive real algebraic numbers, therefore we can consider their (positive) square roots $\sqrt{\zeta_1}, \dots, \sqrt{\zeta_s}$. These are the positive roots of the polynomial $f(t^2)$. Isolating intervals and minimal polynomials (over \mathbb{Q}) of the real algebraic numbers $\sqrt{\zeta_1}, \dots, \sqrt{\zeta_s}$ can be efficiently found by the method of [9], based on factoring the polynomial $f(t^2)$. Let V_1, \dots, V_s be the eigenspaces of the matrix z . For every $i \in \{1, \dots, s\}$, the eigenspace V_i is defined over $K[\sqrt{\zeta_i}]$ and can be efficiently computed as $V_i = \{v \in \mathbb{R}_0^n \mid zv = \zeta_i v\}$. In fact \mathbb{R}_0^n is a direct sum of these subspaces. We define the matrix on subspaces V_i separately: Let $y_i \in \text{Hom}_{\mathbb{R}_0}(V_i, V)$ be the restriction of y to V_i ($i = 1, \dots, s$). Then $y = y_1 \oplus \dots \oplus y_s$. For every $i \in \{1, \dots, s\}$ we set $x_i \in \text{Hom}_{\mathbb{R}_0}(V_i, V)$ as $x_i = \sqrt{\zeta_i}^{-1} y_i$. The argument used in the claim shows that $x = x_1 \oplus \dots \oplus x_s$ is an orthogonal matrix satisfying $(*)$. This finishes the proof of Theorem 4. \square

We remark that the problem of simultaneous conjugacy over the unitary group $\text{U}(\mathbb{C}_0)$ can be treated in a similar way.

4 Decomposition of modules

Let V be a module over the algebra Λ . Recall that V is *decomposable* if $V = W_1 \oplus \dots \oplus W_m$ for some submodules $0 < W_1, \dots, W_m < V$, and *indecomposable* otherwise. V decomposes as a direct sum of indecomposable submodules V_1, \dots, V_r . By the Krull-Schmidt theorem, the isomorphism classes

(counted with multiplicities) of these indecomposable components are unique.

There is a well known correspondence between decompositions of the Λ -module V and decompositions of the identity element of the centralizer algebra $\text{End}_\Lambda(V)$ as direct sums of pairwise orthogonal idempotents. Indeed, assume that $V = W_1 \oplus \dots \oplus W_m$, where $0 < W_1, \dots, W_m \leq V$. For $i = 1, \dots, m$, let π_i denote the projection onto the component W_i with kernel $\sum_{j \neq i} W_j$. It is straightforward to see that the π_i are pairwise orthogonal idempotents in $\text{End}_\Lambda(V)$ with $\sum_{i=1}^m \pi_i = \text{Id}_V$. On the other hand, assume that $\text{Id}_V = \pi_1 + \dots + \pi_m$ be a decomposition of Id_V as a sum of pairwise orthogonal idempotents from $\text{End}_\Lambda(V)$. Then V decomposes as the direct sum of the subspaces $W_1 = \text{im } \pi_1, \dots, W_m = \text{im } \pi_m$. A component W_i of a decomposition of V is indecomposable if and only if the corresponding projection π_i is a primitive idempotent in $\text{End}_\Lambda(V)$.

Thus, finding decompositions of modules is related to the problem of finding orthogonal systems of idempotents in algebras. We show a more direct connection in the case where our module is the regular Λ -module ${}_\Lambda\Lambda$. Assume that Λ is a direct sum of the left ideals L_1, \dots, L_m . Then we can write $1_\Lambda = e_1 + \dots + e_m$ with $e_i \in L_i$. It is straightforward to see that e_1, \dots, e_m are pairwise orthogonal idempotents. On the other hand, if 1_Λ decomposes as a sum of pairwise orthogonal idempotents e_1, \dots, e_m then Λ is a direct sum of the left ideals $L_1 = \Lambda e_1, \dots, L_m = \Lambda e_m$. Again, indecomposable summands of ${}_\Lambda\Lambda$ correspond to primitive idempotents. The (isomorphism classes of) the indecomposable summands of ${}_\Lambda\Lambda$ are called the *principal indecomposable* Λ -modules.

Thus, decomposition of algebras as direct sums of left ideals or, equivalently, finding orthogonal systems of idempotents plays a key role in module decompositions.

4.1 Decomposition of semisimple algebras

In this subsection we recall some known algorithms for decomposing semisimple algebras.

By [7] and [13], a complete system of primitive idempotents in a semisimple algebra Λ over the finite field K can be found in Las Vegas polynomial time. We remark that there are deterministic versions of these algorithms which run in time $(\text{input size} + p)^{O(1)}$, where p is the characteristic of the ground field.

As for the algebraically or real closed case, we summarize the results of Eberly [5], adapted to our needs. The algorithm proposed by Eberly was originally a Las Vegas method. The deterministic version is due to Rónyai, [14]. Let Λ be a semisimple algebra over the algebraic number field K , where K is assumed to be real. Then a complete orthogonal system of primitive idempotents e_1, \dots, e_m of Λ^E can be found by a deterministic polynomial time algorithm. Every idempotents e_i is defined over a number field $K \leq L_i < E$ with

$$\dim_K L_i \leq \begin{cases} \binom{\dim_K \Lambda}{2} & (E = \mathbb{R}_0); \\ \dim_K \Lambda & (E = \mathbb{C}_0). \end{cases}$$

The fields L_i generally differ for different indices i . Note however, that we could take some restrictions. For example, the field of definition of the primitive idempotents e_i can be required to depend only on the simple component Λ^E containing e_i .

4.2 Decomposing arbitrary algebras

In order to find idempotents in an arbitrary K -algebra Λ we use a general procedure that lifts the semisimple part. First of all, we compute the radical $\text{Rad}\Lambda$. Over algebraic number fields, the standard method by Dickson is available. Over finite fields we use the method of [13]. By the Principal Theorem of Wedderburn and Malcev, (c.f. [11]) there exists a subalgebra Δ of Λ isomorphic to the semisimple part $\Lambda/\text{Rad}\Lambda$. In [4], a polynomial time algorithm is presented for finding such a subalgebra, where the ground field K is either a finite field or an algebraic number field. Since every primitive idempotent of Δ is a primitive idempotent in Λ (and a similar statement holds for idempotents in Δ^E), we can apply the methods available for semisimple algebras. We obtain the following results.

Proposition 11 *Let Λ be a K -algebra, where K is a finite field. There is a Las Vegas polynomial time algorithm that finds a complete orthogonal system of primitive idempotents in Λ . The same task can be done by a deterministic method running in time $(\text{input size} + p)^{O(1)}$. \square*

Proposition 12 *Let Λ be a K -algebra, where K is a algebraic number field. Let E be the algebraic or real closure of K . In the latter case we require K to*

be a real field. There is a polynomial time algorithm that finds a complete orthogonal system of primitive idempotents e_1, \dots, e_m in Λ^E . The idempotents e_i are defined over the (real) algebraic number fields L_i . \square

Now we can prove Theorems 5 and 6 as follows. Let V be a Λ -module. We can compute the centralizer $\text{End}_\Lambda(V)$ by solving a system of linear equations. We find a complete orthogonal system π_1, \dots, π_m of primitive idempotents in $\text{End}_\Lambda(V)$ by the method of Proposition 11 or 12, respectively, and compute the submodules $V_i = \pi_i(V)$ ($i = 1, \dots, m$). This gives a decomposition of V as the direct sum of indecomposable modules. We have finished the proof of Theorems 5 and 6. \square

References

- [1] BABAI, L., AND RÓNYAI, L. Computing irreducible representations of finite groups. *Mathematics of Computation* 55, 192 (1990), 705–722.
- [2] CHISTOV, A. L., AND KARPINSKI, M. Polynomial time decomposition of modules over algebras and its application. Tech. Rep. 85152-CS, Institut für Informatik der Universität Bonn, 1996.
- [3] COHEN, A. M., IVANYOS, G., AND WALES, D. B. Finding the radical of an algebra of linear transformations. *J. Pure and Applied Algebra* (1997), to appear. (Spec. issue: Proc. 4th Int. Symp. on Effective Methods in Algebraic Geometry).
- [4] DE GRAAF, W. A., IVANYOS, G., KÜRONYA, A., AND RÓNYAI, L. Computing Levi decompositions. *Applicable Algebra in Engineering, Communication and Computing* (to appear).
- [5] EBERLY, W. M. Decompositions of algebras over \mathbf{R} and \mathbf{C} . *Computational Complexity* 1 (1991), 179–206.
- [6] EDMONDS, J. System of distinct representatives and linear algebra. *Journal of Research of the National Bureau of Standards* 718, 4 (1967), 241–245.

- [7] FRIEDL, K., AND RÓNYAI, L. Polynomial time solution of some problems in computational algebra. In *Proc. 17th ACM STOC* (1985), pp. 153–162.
- [8] HOLT, D. F., AND REES, S. Testing modules for irreducibility. *J. Austral. Math. Soc. Ser. A* 57 (1994), 1–16.
- [9] LANDAU, S. Factoring polynomials over algebraic number fields. *SIAM J. on Computing* 14 (1985), 184–195.
- [10] PARKER, R. A. The computer calculation of modular characters (the Meat-Axe). In *Computational Group Theory* (1984), Academic Press, pp. 267–274.
- [11] PIERCE, R. S. *Associative Algebras*. Springer-Verlag, 1982.
- [12] RÓNYAI, L. Zero divisors in quaternion algebras. *Journal of Algorithms* 9 (1988), 494–506.
- [13] RÓNYAI, L. Computing the structure of finite algebras. *J. Symbolic Computation* 9 (1990), 355–373.
- [14] RÓNYAI, L. A deterministic method for computing splitting elements in semisimple algebras over \mathbb{Q} . *Journal of Algorithms* 16 (1994), 24–32.
- [15] SCHWARTZ, J. T. Fast probabilistic algorithms for verification of polynomial identities. *Journal of ACM* 27 (1980), 701–717.
- [16] ZIPPEL, R. E. Probabilistic algorithms for sparse polynomials. In *Proc. EUROSAM '79* (1979), vol. 72 of *Lect. Notes in Comp. Sci.*, Springer, pp. 216–226.