A $\frac{5}{2}n^2$–Lower Bound for the Multiplicative Complexity of $n \times n$–Matrix Multiplication

Markus Bläser

Institut für Theoretische Informatik, Med. Universität zu Lübeck
Wallstr. 40, 23560 Lübeck, Germany
blaeser@tcs.mu-luebeck.de

Abstract. We prove a lower bound of $\frac{5}{2}n^2 - 3n$ for the multiplicative complexity of $n \times n$–matrix multiplication over arbitrary fields. More general, we show that for any finite dimensional semisimple algebra A with unity, the multiplicative complexity of the multiplication in A is bounded from below by $\frac{5}{2}\dim A - 3(n_1 + \cdots + n_t)$ if the decomposition of $A \cong A_1 \times \cdots \times A_t$ into simple algebras $A_\tau \cong D_\tau{}^{n_\tau \times n_\tau}$ contains only noncommutative factors, that is, the division algebra D_τ is noncommutative or $n_\tau \geq 2$.

1 Introduction

One of the leading problems in algebraic complexity theory is the determination of good (lower as well as upper) bounds for the multiplicative complexity of $n \times n$–matrix multiplication. Loosely speaking, the problem is the following: given $n \times n$–matrices $X = (X_{i,j})$ and $Y = (Y_{i,j})$ with indeterminates $X_{i,j}$ and $Y_{i,j}$ over some ground field k, how many *essential* multiplications and divisions are needed to compute the entries of the product XY? Here, "essential" means that additions, subtractions, and scalar multiplications are free of costs. According to Strassen [20], we may reformulate the problem over infinite fields as follows: the multiplicative complexity of $n \times n$–matrix multiplication is the smallest number ℓ of products

$$p_\lambda = u_\lambda(X_{i,j}, Y_{i,j}) \cdot v_\lambda(X_{i,j}, Y_{i,j})$$

with linear forms u_λ and v_λ in the $X_{i,j}$ and $Y_{i,j}$ such that each entry of XY is contained in the linear span of $p_1, \ldots, p_\ell$, i.e.,

$$\sum_{\nu=1}^{n} X_{i,\nu} Y_{\nu,j} \in \operatorname{lin}\{p_1, \ldots, p_\ell\} \qquad \text{for } 1 \leq i, j \leq n.$$

In other words, we may restrict our attention to computations that contain only "normalized" multiplications and no divisions. (Since we are considering lower bounds in this work, the above restriction to infinite fields does not impose any problems: lower bounds over a field K also hold over any subfield $k \subset K$.) For a modern and comprehensive introduction to algebraic complexity theory, we refer to [9].

A related quantity is the bilinear complexity (or rank) of $n \times n$–matrix multiplication. Here the products $p_\lambda = u_\lambda(X_{i,j}) \cdot v_\lambda(Y_{i,j})$ are bilinear products, that is, products of linear forms u_λ in the $X_{i,j}$ and linear forms v_λ in the $Y_{i,j}$. (Note that the entries of XY are bilinear forms.) The concept of bilinear complexity has been utilized with great success in the design of asymptotically fast matrix multiplication algorithms, see for example [19, 2, 18, 22, 10]. Obviously, the multiplicative complexity is a lower bound for the bilinear complexity and it is not hard to see that twice the multiplicative complexity is an upper bound for the bilinear complexity (see e.g. [9, Eq. 14.8]). Therefore, we usually want to have upper bounds for the bilinear complexity and lower bounds for the multiplicative complexity. While the difference between multiplicative and bilinear complexity seems to be minor at a first glance, it is much harder to cope with the multiplicative complexity when dealing with lower bounds. One reason among others is the fact that the bilinear complexity of a tensor of a bilinear map (see below for a definition) is invariant under permutations whereas the multiplicative complexity might not, see also [9, Chap. 14.2] for a further discussion.

1.1 Model of Computation

Of course, we can define multiplicative complexity not only for the multiplication of $n \times n$–matrices (which is a bilinear map $k^{n\times n} \times k^{n\times n} \to k^{n\times n}$) but also for arbitrary bilinear maps. When considering lower bounds, it is often more convenient to use a coordinate-free definition of multiplicative complexity, see e.g. [9, Chap. 14.1]. In the following, if V is a vector space, let V^* denote the dual space of V, i.e., the vector space of all linear forms on V.

Definition 1. *Let k be a field, U, V, and W finite dimensional vector spaces over k, and $\phi : U \times V \to W$ a bilinear map.*

1. *A sequence $\beta = (f_1, g_1, w_1, \dots, f_\ell, g_\ell, w_\ell)$ with $f_\lambda, g_\lambda \in (U \times V)^*$ and $w_\lambda \in W$ is called a quadratic computation for ϕ over k of length ℓ if*

$$\phi(u, v) = \sum_{\lambda=1}^{\ell} f_\lambda(u, v) g_\lambda(u, v) w_\lambda \qquad \text{for all } u \in U, v \in V.$$

2. *The length of a shortest quadratic computation for ϕ is called the multiplicative complexity of ϕ and is denoted by $C(\phi)$.*
3. *If A is a finite dimensional associative k-algebra with unity, then the multiplicative complexity of A is defined as the multiplicative complexity of the multiplication map of A, which is a bilinear map $A \times A \to A$, and is denoted by $C(A)$.*

If we want to emphasize the underlying ground field k, we will sometimes write $C_k(\phi)$ and $C_k(A)$ instead of $C(\phi)$ and $C(A)$, respectively. Using the language from above, the multiplicative complexity of $n \times n$–matrix multiplication is denoted by $C(k^{n\times n})$.

If we require that $f_\lambda \in U^*$ and $g_\lambda \in V^*$ in the above Definition 1, we get *bilinear computations* and *bilinear complexity* (also called *rank*). We denote the bilinear complexity of a bilinear map ϕ by $R(\phi)$ or $R_k(\phi)$ and the bilinear complexity of an associative algebra A by $R(A)$ or $R_k(A)$. For any bilinear map ϕ, we have

$$C(\phi) \leq R(\phi) \leq 2 \cdot C(\phi)\,. \tag{1}$$

Except for trivial cases, the second inequality is always strict, see [14].

1.2 Previous Results

In 1978, Brockett and Dobkin [7] proved the bound $R(k^{n\times n}) \geq 2n^2-1$ and in the same year, Lafon and Winograd [15] extended this result to the multiplicative complexity. Three years later, Alder and Strassen [1] unified most of the lower bounds known at that time, including the last one, in a single theorem: for any finite dimensional associative k-algebra A

$$C(A) \geq 2 \dim A - t\,, \tag{2}$$

where t is the number of maximal twosided ideals in A. Recently, some progress has been made in the case of matrix multiplication, namely $C(k^{n\times n}) \geq 2n^2+n-3$ (see [4]). Bshouty [8] obtained the bound $R_{GF(2)}(GF(2)^{n\times n}) \geq \frac{5}{2}n^2 - o(n^2)$ for the special case $k = GF(2)$ using methods from coding theory. He claims that this bound also holds for the multiplicative complexity (over $GF(2)$) but does not give a proof. Finally, in [3] we proved the lower bound

$$R(k^{n\times n}) \geq \tfrac{5}{2}n^2 - 3n\,. \tag{3}$$

for arbitrary fields k.

1.3 New Results

As our first main result, we show that (3) also holds for the multiplicative complexity. The following theorem is proven in Section 4.

Theorem 1. *For any field k, $C(k^{n\times n}) \geq \frac{5}{2}n^2 - 3n$.*

One of the main ingredients of its proof is Lemma 5 which is essentially a new combination of a result by Ja'Ja' [14] on the relation between multiplicative and bilinear complexity, Strassen 3-slice tensor technique [21], and the substitution method [16]. We prove Lemma 5 in Section 3. Before doing so, we introduce the so-called "tensorial notion" in the next Section 2 and present a (well-known) alternative characterization of multiplicative complexity, which is more suited for our purposes.

The bound of Theorem 1 is a special case of the following lower bound (which we show in Section 5).

Theorem 2. *Let $A \cong A_1 \times \cdots \times A_t$ be a semisimple algebra over an arbitrary field k with $A_\tau = D_\tau{}^{n_\tau \times n_\tau}$ for all τ, where D_τ is a k-division algebra. Assume that each factor A_τ is noncommutative, that is, $n_\tau \geq 2$ or D_τ is noncommutative. Moreover, let $n = n_1 + \cdots + n_t$. Then $C(A) \geq \frac{5}{2} \dim A - 3n$.*

Our new bound of Theorem 2 is the first lower bound over arbitrary fields for the multiplicative complexity of a semisimple algebra—in particular of the important algebra $k^{n\times n}$—significantly above the Alder–Strassen bound (2). In [5] (see also the forthcoming [6]), we obtained the same bound for the easier case of the bilinear complexity. While in this case, our bounds can also be extended to arbitrary finite dimensional algebras A provided that the (semisimple) quotient algebra $A/\operatorname{rad} A$ satisfies the assumptions of Theorem 2, we provide an example in Section 6 which shows that our methods cannot yield this generalization for the multiplicative complexity (without any extra considerations). This again gives evidence for the intricate nature of multiplicative complexity.

2 Characterizations of Multiplicative Complexity

In the previous section, we have introduced the multiplicative complexity of a bilinear map in terms of computations. A second useful characterization of multiplicative complexity is the so-called “tensorial notion” (see [9, Chap. 14.4] for the bilinear complexity). With a bilinear map $\phi : U \times V \to W$, we may associate a *coordinate tensor* (or tensor for short) which is basically a “three-dimensional matrix”: we fix bases $u_1, \ldots, u_m$ of U, $v_1, \ldots, v_n$ of V, and $w_1, \ldots, w_p$ of W. There are unique scalars $t_{\mu,\nu,\rho} \in k$ such that

$$\phi(u_\mu, v_\nu) = \sum_{\rho=1}^{p} t_{\mu,\nu,\rho} w_\rho \qquad \text{for all } 1 \leq \mu \leq m,\ 1 \leq \nu \leq n. \tag{4}$$

Then $t = (t_{\mu,\nu,\rho}) \in k^{m\times n\times p}$ is the tensor of ϕ (with respect to the chosen bases). On the other hand, any given tensor also defines a bilinear map after choosing bases. We define the multiplicative complexity of the tensor t by $C(t) := C(\phi)$. In the same way, the bilinear complexity of t is $R(t) := R(\phi)$. (This is in both cases well-defined, since the multiplicative resp. bilinear complexity is robust with respect to invertible linear transformations, i.e., to changes of bases.)

If ϕ is the multiplication in an algebra A, then we may instantiate the above three bases with one and the same basis. In this case, the tensor consists of the structural constants (see [11] for a definition) of the algebra A (with respect to the chosen basis).

With each tensor $t = (t_{\mu,\nu,\rho})$, we may associate three sets of matrices, the *slices* of t. The matrices $Q_\mu = (t_{\mu,\nu,\rho})_{1\leq\nu\leq n, 1\leq\rho\leq p} \in k^{n\times p}$ with $1 \leq \mu \leq m$ are called the 1-slices of t, the matrices $S_\nu = (t_{\mu,\nu,\rho})_{1\leq\mu\leq m, 1\leq\rho\leq p} \in k^{m\times p}$ with $1 \leq \nu \leq n$ the 2-slices, and finally $T_\rho = (t_{\mu,\nu,\rho})_{1\leq\mu\leq m, 1\leq\nu\leq n} \in k^{m\times n}$ with $1 \leq \rho \leq p$ are called the 3-slices of t. When dealing with bilinear complexity, it makes no difference which of the three sets of slices we consider. In the case of multiplicative complexity, however, the 3-slices play a distinguished role. (This is one reason why proving lower bounds for the multiplicative complexity is hard.)

Lemma 1. *Let k be a field and t be a tensor with 3-slices $T_1, \dots, T_p \in k^{m\times n}$. Then $C(t) \le \ell$ if and only if there are (column) vectors $u_\lambda, v_\lambda \in k^{m+n}$ for $1 \le \lambda \le \ell$ such that with $P_\lambda := u_\lambda \cdot v_\lambda^\top \in k^{(m+n)\times(m+n)}$*

$$\begin{pmatrix} 0 & T_1 \\ T_1^\top & 0 \end{pmatrix}, \dots, \begin{pmatrix} 0 & T_p \\ T_p^\top & 0 \end{pmatrix} \in \mathrm{lin}\{P_1 + P_1^\top, \dots, P_\ell + P_\ell^\top\}\,. \tag{5}$$

Here, $T^\top$ denotes the transpose of a matrix T and $\mathrm{lin}\{\dots\}$ denotes the linear span. A proof of this lemma is straight forward. (One possibility is to follow the lines of the proof of Theorem 3.2 in [14].) The rank one matrices P_λ correspond to the products of a quadratic computation. By transposing, we identify the product xy of two indeterminates with yx.

If $T_1, \dots, T_p$ are the 3-slices of a tensor t, we will occasionally also write $C(T_1, \dots, T_p)$ instead of $C(t)$ and $R(T_1, \dots, T_p)$ instead of $R(t)$. By multiplying (5) with

$$\begin{pmatrix} X & 0 \\ 0 & Y^\top \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} X^\top & 0 \\ 0 & Y \end{pmatrix}$$

from the left and right, respectively, it follows from Lemma 1 that if $X \in k^{m\times m}$ and $Y \in k^{n\times n}$ are invertible matrices, then

$$C(T_1, \dots, T_p) = C(X \cdot T_1 \cdot Y, \dots, X \cdot T_p \cdot Y)\,. \tag{6}$$

This multiplication of the slices with X and Y corresponds to a change of the bases $u_1, \dots, u_m$ and $v_1, \dots, v_n$ in (4).

3 Lower Bounds

In the present section, we prove our main lemma (Lemma 5). Its proof is done by a combination of the so-called substitution method (first used for proving lower bounds in algebraic complexity theory by Pan [16]), Strassen's lower bound for the (border) rank of a 3-slice tensor [21], and a result by Ja'Ja' [14] which relates multiplicative and bilinear complexity in a more sophisticated way than (1).

We first state the results of Strassen and Ja'Ja'. For two elements b, c of an associative algebra, let $[b, c] := bc - cb$ denote their *Lie product* (or *commutator*). For a matrix $M \in k^{N\times N}$, let $\mathrm{rk}\, M$ denote its (usual) rank. We denote the identity matrix of $k^{N\times N}$ by I_N.

Lemma 2 (Strassen). *Let t be a tensor with 3-slices $I_N, B, C \in k^{N\times N}$ over some field k. Then $R(t) \ge N + \frac{1}{2}\mathrm{rk}[B, C]$.*

For a proof, see [21, Thm. 4.1] or [9, Thm. 19.12]. Strassen actually proves the above lemma for the so-called border rank of t (see e.g. [9, Chap. 15.4] for a definition) which is a lower bound for the bilinear complexity.

Lemma 3 (Ja'Ja'). *Let t be a tensor with 3-slices $T_1,\dots,T_p \in k^{N\times N}$ over some field k. Then*

$$C(t) \geq \tfrac{1}{2} R\Big(\begin{pmatrix} T_1 & 0 \\ 0 & T_1^\top \end{pmatrix}, \dots, \begin{pmatrix} T_p & 0 \\ 0 & T_p^\top \end{pmatrix} \Big).$$

For a proof, see [14, Thm. 3.4].

Combining these two lemmata, we obtain a lower bound for the multiplicative complexity of a 3-slice tensor.

Lemma 4. *Let k be a field. Let t be a tensor with 3-slices $I_N, B, C \in k^{N\times N}$. Then $C(t) \geq N + \frac{1}{2}\operatorname{rk}[B,C]$.*

Proof. By the previous Lemmata 3 and 2

$$\begin{aligned} C(t) &\geq \tfrac{1}{2} R\Big(I_{2N}, \begin{pmatrix} B & 0 \\ 0 & B^\top \end{pmatrix}, \begin{pmatrix} C & 0 \\ 0 & C^\top \end{pmatrix} \Big) \\ &\geq \tfrac{1}{2}\Big(2N + \tfrac{1}{2} \operatorname{rk} \Big[\begin{pmatrix} B & 0 \\ 0 & B^\top \end{pmatrix}, \begin{pmatrix} C & 0 \\ 0 & C^\top \end{pmatrix} \Big] \Big). \end{aligned} \tag{7}$$

Since $B^\top C^\top - C^\top B^\top = -[B,C]^\top$,

$$\operatorname{rk} \Big[\begin{pmatrix} B & 0 \\ 0 & B^\top \end{pmatrix}, \begin{pmatrix} C & 0 \\ 0 & C^\top \end{pmatrix} \Big] = \operatorname{rk} \begin{pmatrix} BC - CB & 0 \\ 0 & B^\top C^\top - C^\top B^\top \end{pmatrix} = 2 \operatorname{rk}[B,C].$$

Thus, the right-hand side of (7) equals $N + \frac{1}{2}\operatorname{rk}[B,C]$. □

We now come to the proof of our main lemma. We combine the substitution method (here manifesting itself in the Steinitz exchange) with the previous Lemma 4.

Lemma 5. *Let t be a tensor with linearly independent 3-slices $T_1,\dots,T_p \in k^{N\times N}$ over some field k. Assume there are integers s and q such that for any basis $U_1,\dots,U_p$ of $\operatorname{lin}\{T_1,\dots,T_p\}$ there are indices $i_1,\dots,i_s$ and $j_1,\dots,j_q$ with the following properties: the linear span of $U_{i_1},\dots,U_{i_s}$ contains an invertible matrix E and the linear span of $U_{j_1},\dots,U_{j_q}$ contains matrices B and C with $\operatorname{rk}[B,C] = N$. Then $C(t) \geq p - s - q + \frac{3}{2}N$.*

Proof. Let $\ell := C(t)$. By Lemma 1, there are matrices $P_1,\dots,P_\ell$ of rank one and $S_\rho \in \operatorname{lin}\{P_1,\dots,P_\ell\}$ such that

$$\begin{pmatrix} 0 & T_\rho \\ T_\rho^\top & 0 \end{pmatrix} = S_\rho + S_\rho^\top \qquad \text{for } 1 \leq \rho \leq p. \tag{8}$$

The matrices $S_1,\dots,S_p$ are linearly independent, since otherwise we would obtain a linear dependence of the 3-slices $T_1,\dots,T_p$ from (8), a contradiction. We now exploit the Steinitz exchange in a rather explicit way: write $S_\rho = \sum_{\lambda=1}^{\ell} \xi_{\rho,\lambda} P_\lambda$ with scalars $\xi_{\rho,\lambda} \in k$ for $1 \leq \rho \leq p$. Let $X = (\xi_{\rho,\lambda}) \in k^{p\times \ell}$.

The matrix X has full rank p, since the S_ρ are linearly independent. By permuting the P_λ, we may assume that the matrix X' consisting of the first p columns of X is invertible. Let Y' be its inverse. We may augment Y' to a matrix $Y = (\eta_{\rho,\lambda}) \in k^{p\times \ell}$ such that

$$\underbrace{\sum_{i=1}^{p} \eta_{\rho,i} S_i}_{=:\, M_\rho} = P_\rho + \sum_{\lambda=p+1}^{\ell} \eta_{\rho,\lambda} P_\lambda \qquad \text{for } 1 \le \rho \le p. \tag{9}$$

The above defined matrices $M_1, \dots, M_p$ are linearly independent and their linear span equals $\operatorname{lin}\{S_1, \dots, S_p\}$. By virtue of (8),

$$\begin{pmatrix} 0 & T_\rho \\ T_\rho^\top & 0 \end{pmatrix} \in \operatorname{lin}\{M_1 + M_1^\top, \dots, M_p + M_p^\top\} \qquad \text{for } 1 \le \rho \le p. \tag{10}$$

Let L_ρ be the $N \times N$–matrix defined by

$$\begin{pmatrix} 0 & L_\rho \\ L_\rho^\top & 0 \end{pmatrix} = M_\rho + M_\rho^\top .$$

($M_\rho + M_\rho^\top$ is of the above form by (10) and the linear independence of $T_1, \dots, T_p$.) By (10), $\operatorname{lin}\{L_1, \dots, L_p\}$ equals $\operatorname{lin}\{T_1, \dots, T_p\}$. By the assumption of the lemma, there are indices $i_1, \dots, i_s$ such that $\operatorname{lin}\{L_{i_1}, \dots, L_{i_s}\}$ contains an invertible matrix E. Exploiting (6), we may replace L_ρ with $E^{-1}L_\rho$ for $1 \le \rho \le p$. Thereafter, $I_N \in \operatorname{lin}\{L_{i_1}, \dots, L_{i_s}\}$. Again by assumption, there are indices $j_1, \dots, j_q$ such that $\operatorname{lin}\{L_{j_1}, \dots, L_{j_q}\}$ contains matrices B and C with $\operatorname{rk}[B, C] = N$.

By (9), it follows that

$$\begin{aligned} \begin{pmatrix} 0 & I_N \\ I_N^\top & 0 \end{pmatrix}, \begin{pmatrix} 0 & B \\ B^\top & 0 \end{pmatrix}, \begin{pmatrix} 0 & C \\ C^\top & 0 \end{pmatrix} \in \operatorname{lin}\{ & P_{i_1} + P_{i_1}^\top, \dots, P_{i_s} + P_{i_s}^\top, \\ & P_{j_1} + P_{j_1}^\top, \dots, P_{j_q} + P_{j_q}^\top, \\ & P_{p+1} + P_{p+1}^\top, \dots, P_\ell + P_\ell^\top\}. \end{aligned}$$

Thus, $C(I_N, B, C) \le \ell - p + s + q$ yielding $C(t) \ge p - s - q + C(I_N, B, C)$. By Lemma 4, $C(I_N, B, C) \ge \frac{3}{2}N$. □

4 Matrix Multiplication

As the first and most important example, we apply Lemma 5 to the algebra $k^{n\times n}$. Our aim is to utilize the following two lemmata which are proven in [3, Sect. 4].

Lemma 6. *Let k be an infinite field and let V be a subspace of $k^{n\times n}$ that contains an invertible matrix. Then for any basis $v_1, \dots, v_d$ of V there are $s \le n$ and indices $i_1, \dots, i_s$ such that already the linear span of $v_{i_1}, \dots, v_{i_s}$ contains an invertible matrix.*

Lemma 7. *Let k be an infinite field, let $n \geq 2$, and let $v_1, \ldots, v_p$ be a basis of $k^{n\times n}$, where $p = n^2$. Then there are $q \leq 2n$, indices $j_1, \ldots, j_q$, and $b, c \in \operatorname{lin}\{v_{j_1}, \ldots, v_{j_q}\}$ such that $[b, c]$ is invertible.*

To exploit the above two lemmata, we have to relate the structure of the algebra $k^{n\times n}$ to the structure of the 3-slices of the coordinate tensor of $k^{n\times n}$.

For the moment, consider an arbitrary associative algebra A of dimension p. For an element $x \in A$, let ℓ_x and r_x denote the vector space endomorphisms defined by $y \mapsto xy$ and $y \mapsto yx$ for all $y \in A$, respectively. Let $a_1, \ldots, a_p$ be a basis of A and t be the corresponding coordinate tensor. From (4) (where each of the three bases is instantiated with $a_1, \ldots, a_p$) it is clear that the ρth 1-slice of t is the matrix of the left multiplication ℓ_{a_ρ} (with respect to $a_1, \ldots, a_p$). Thus, the homomorphism that maps a_ρ to the ρth 1-slice of t for each ρ is a faithful representation of A. Therefore, the subalgebra of $k^{p\times p}$ generated by the 1-slices of t is isomorphic to A. In the same way, the subalgebra of $k^{p\times p}$ generated by the 2-slices is isomorphic to the opposite algebra A^o of A (see [11] for a definition). The question how the 3-slices of t are related to the structure of A is a more subtle question.

In the case of the algebra $k^{n\times n}$, we are lucky. It is well known that the structure of the coordinate tensor t of $k^{n\times n}$ is invariant under permutations, see for example [9, Eq. 14.21]. It follows that if $Q_1, \ldots, Q_p$ denote the 1-slices of t and $T_1, \ldots, T_p$ denote the 3-slices of t (where $p = n^2$), then there are invertible matrices $X \in k^{p\times p}$ and $Y \in k^{p\times p}$ (even permutation matrices) such that

$$\operatorname{lin}\{Q_1, \ldots, Q_p\} = \operatorname{lin}\{X \cdot T_1 \cdot Y, \ldots, X \cdot T_p \cdot Y\}\,.$$

By (6), we may replace T_ρ by $X \cdot T_\rho \cdot Y$ for $1 \leq \rho \leq p$. Thereafter, the subalgebra of $k^{p\times p}$ generated by $T_1, \ldots, T_p$ is isomorphic to the algebra $k^{n\times n}$. By the above Lemmata 6 and 7, $T_1, \ldots, T_p$ fulfil the assumptions of Lemma 5 with $s = n$ and $q = 2n$. (The restriction that k is infinite imposes no problem here, since $C_k(k^{n\times n}) \geq C_K(K^{n\times n})$ for any extension field $K \supset k$.) Now Lemma 5 yields $C(T_1, \ldots, T_p) \geq n^2 - n - 2n + \frac{3}{2}n^2$. This completes the proof of Theorem 1.

5 Further Applications: Semisimple Algebras

In the present section, we generalize the results of the preceding section to semisimple algebras.

As seen above, we have to relate the 3-slices of the coordinate tensor of an algebra with the structure of that algebra. We start with the case of a simple k-algebra A. By Wedderburn's Structure Theorem (see [11]), A is isomorphic to an algebra $D^{n\times n} \cong D \otimes k^{n\times n}$ for some positive integer n and some k-division algebra D. Let $p := \dim A = n^2 \cdot \dim D$. Let $a_1, \ldots, a_p$ be a basis of A and let $a_1^*, \ldots, a_p^*$ denote its dual basis. Let $\ell_{a_i}^*$ denote the dual of the left multiplication with a_i, that is, the linear map $A^* \to A^*$ defined by $b \mapsto b \circ \ell_{a_i}$. Applying a_ρ^* to (4) yields

$$a_\rho^*(\ell_{a_\mu}(a_\nu)) = t_{\mu,\nu,\rho}\,,$$

hence

$$\ell^*_{a_\mu}(a^*_\rho) = \sum_{\nu=1}^{p} t_{\mu,\nu,\rho} a^*_\nu \qquad \text{for all } 1 \le \mu, \rho \le p. \tag{11}$$

De Groote [12, Prop. 1.1] shows that for any simple algebra A there is a vector space isomorphism $S : A \to A^*$ fulfilling

$$\ell^*_x = S \circ r_x \circ S^{-1} \qquad \text{for all } x \in A.$$

Substituting this into (11), we obtain

$$S^{-1}(a^*_\rho) \cdot a_\mu = \sum_{\nu=1}^{p} t_{\mu,\nu,\rho} S^{-1}(a^*_\nu) \qquad \text{for all } 1 \le \mu, \rho \le p.$$

Thus, the ρth 3-slice of the tensor of A with respect to the basis $a_1, \ldots, a_p$ is the matrix of the homomorphism $\ell_{S^{-1}(a^*_\rho)} : A \to A$ with respect to the two bases $a_1, \ldots, a_p$ and $S^{-1}(a^*_1), \ldots, S^{-1}(a^*_p)$. Hence, there are invertible matrices $X, Y \in k^{p \times p}$ such that the subalgebra of $k^{p \times p}$ generated by the matrices $X \cdot T_1 \cdot Y, \ldots, X \cdot T_p \cdot Y$ is isomorphic to A.

Next, we consider the case of a semisimple algebra. By Wedderburn's Structure Theorem, any semisimple algebra A is isomorphic to a direct product of simple algebras $A_1 \times \cdots \times A_t$ where each $A_\tau = {D_\tau}^{n_\tau \times n_\tau}$ for some k-division algebra D_τ. If we choose a basis with respect to this decomposition of A, then the corresponding coordinate tensor is a direct sum of the tensors of $A_1, \ldots, A_t$. By applying the above considerations for the simple case separately to each A_τ, we conclude that there are invertible matrices $X, Y \in k^{p \times p}$ such that the subalgebra of $k^{p \times p}$ spanned by $X \cdot T_1 \cdot Y, \ldots, X \cdot T_p \cdot Y$ is isomorphic to A.

The following analogue of Lemma 6 and Lemma 7 is proven in [5]:

Lemma 8. *Let $A \cong A_1 \times \cdots \times A_t$ be a semisimple algebra over an infinite field k with $A_\tau = {D_\tau}^{n_\tau \times n_\tau}$ for all τ, where D_τ is a k-division algebra. Assume that each factor A_τ is noncommutative, that is, $n_\tau \ge 2$ or D_τ is noncommutative. Moreover, let $n = n_1 + \cdots + n_t$ and $v_1, \ldots, v_p$ be a basis of A.*

1. *There are $s \le n$ and indices $i_1, \ldots, i_s$ such that $\mathrm{lin}\{v_{i_1}, \ldots, v_{i_s}\}$ contains an invertible element.*
2. *There are $q \le 2n$, indices $j_1, \ldots, j_q$, and $b, c \in \mathrm{lin}\{v_{j_1}, \ldots, v_{j_q}\}$ such that $[b, c]$ is invertible.*

By the above Lemma 8, the 3-slices $T_1, \ldots, T_p$ fulfil the assertion of Lemma 5 with $s = n$ and $q = 2n$ where $n = n_1 + \cdots + n_t$. (Again, the restriction that k is infinite does not impose any problems here, we can switch over from k to $k(x)$ with some extra indeterminate x. This does not have any relevant impact on the structure of A.) Altogether, this proves Theorem 2.

6 A Limiting Example

While for the bilinear complexity, our bounds can also be extended to arbitrary finite dimensional algebras A provided that the (semisimple) quotient algebra $A/\operatorname{rad} A$ fulfil the assumptions of Theorem 2 (see [5] and the forthcoming [6]), we here construct an example that satisfies these assumptions but for which our method fails in the case of multiplicative complexity. Of course, this does not mean that our method cannot be applied to arbitrary associative algebras, we just have to examine the 3-slices of the algebra explicitly.

Let $X_1, \dots, X_n$ be indeterminates over some field k. Furthermore, let I denote the ideal generated by all monomials of total degree two. Consider the algebra $A = k[X_1, \dots, X_n]/I$. We have $X_i \cdot X_j = 0$ in A for all i, j. With respect to the basis $1, X_1, \dots, X_n$, the tensor t_A of A looks as follows:

$$t_A = \begin{pmatrix} 1 & 2 & 3 & \cdots & n+1 \\ 2 & & & & \\ 3 & & & & \\ \vdots & & & & \\ n+1 & & & & \end{pmatrix}$$

Above, a ρ in position (μ, ν) means that the ρth 3-slice has the entry one in position (μ, ν). Unspecified entries are zero.

The algebra A is commutative, so $A/\operatorname{rad} A$ does not fulfil the assumptions of Theorem 2. Instead, consider $A' = D \otimes A$ for some noncommutative central division algebra D. We have $A'/\operatorname{rad} A' = D$, thus $A'/\operatorname{rad} A'$ satisfies the assumptions of Theorem 2. However, any matrix in the linear span of the 3-slices of the tensor of A' (which we obtain from t_A by substituting each one in t_A by the tensor of D) has rank at most $2 \dim D$. Consequently, the Lie product of any two such matrices has rank at most $4 \dim D$. Therefore, if n is large, we are not able to obtain the additional $\frac{1}{2} \dim A'$ that we achieved in the bound of Theorem 2.

References

1. A. Alder and V. Strassen. On the algorithmic complexity of associative algebras. *Theoret. Comput. Sci.*, 15:201–211, 1981.
2. Dario Bini, Milvio Capovani, Grazia Lotti, and Francesco Romani. $O(n^{2.7799})$ complexity for matrix multiplication. *Inf. Proc. Letters*, 8:234–235, 1979.
3. Markus Bläser. A $\frac{5}{2}n^2$–lower bound for the rank of $n \times n$–matrix multiplication over arbitrary fields. In *Proc. 40th Ann. IEEE Symp. on Found. Comput. Sci. (FOCS)*, pages 45–50, 1999.
4. Markus Bläser. Lower bounds for the multiplicative complexity of matrix multiplication. *Comput. Complexity*, 8:203–226, 1999.
5. Markus Bläser. *Untere Schranken für den Rang assoziativer Algebren.* Dissertation, Universität Bonn, 1999.
6. Markus Bläser. Lower bounds for the bilinear complexity of associative algebras. *Comput. Complexity*, to appear.

7. Roger W. Brockett and David Dobkin. On the optimal evaluation of a set of bilinear forms. *Lin. Alg. Appl.*, 19:207–235, 1978.
8. Nader H. Bshouty. A lower bound for matrix multiplication. *SIAM J. Comput.*, 18:759–765, 1989.
9. Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi. *Algebraic Complexity Theory.* Springer, 1997.
10. Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progression. *J. Symbolic Comput.*, 9:251–280, 1990.
11. Yurij A. Drozd and Vladimir V. Kirichenko. *Finite Dimensional Algebras.* Springer, 1994.
12. Hans F. de Groote. Characterization of division algebras of minimal rank and the structure of their algorithm varieties. *SIAM J. Comput.*, 12:101–117, 1983.
13. Hans F. de Groote. *Lectures on the Complexity of Bilinear Problems.* LNCS 245. Springer, 1986.
14. Joseph Ja'Ja'. On the complexity of bilinear forms with commutativity. *SIAM J. Comput.*, 9:717–738, 1980.
15. Jean-Claude Lafon and Shmuel Winograd. A lower bound for the multiplicative complexity of the product of two matrices. Technical report, Centre de Calcul de L'Esplanade, U.E.R. de Mathematique, Univ. Louis Pasteur, Strasbourg, France, 1978.
16. Victor Yu. Pan. Methods for computing values of polynomials. *Russ. Math. Surv.*, 21:105–136, 1966.
17. Richard S. Pierce. *Associative Algebras.* Springer, 1982.
18. Arnold Schönhage. Partial and total matrix multiplication. *SIAM J. Comput.*, 10:434–455, 1981.
19. Volker Strassen. Gaussian elimination is not optimal. *Num. Math.*, 13:354–356, 1969.
20. Volker Strassen. Vermeidung von Divisionen. *J. Reine Angew. Math.*, 264:184–202, 1973.
21. Volker Strassen. Rank and optimal computation of generic tensors. *Lin. Alg. Appl.*, 52/53:645–685, 1983.
22. Volker Strassen. Relative bilinear complexity and matrix multiplication. *J. Reine Angew. Math.*, 375/376:406–443, 1987.
23. Volker Strassen. Algebraic complexity theory. In J. van Leeuven, editor, *Handbook of Theoretical Computer Science Vol. A*, pages 634–672. Elsevier Science Publishers B.V., 1990.