

Fast matrix multiplication using coherent configurations

Henry Cohn*

Christopher Umans†

Abstract

We introduce a relaxation of the notion of tensor rank, called *s-rank*, and show that upper bounds on the *s-rank* of the matrix multiplication tensor imply upper bounds on the ordinary rank. In particular, if the “*s-rank exponent of matrix multiplication*” equals 2, then $\omega = 2$. This connection between the *s-rank* exponent and the ordinary exponent enables us to significantly generalize the group-theoretic approach of Cohn and Umans, from group algebras to general algebras. Embedding matrix multiplication into general algebra multiplication yields bounds on *s-rank* (not ordinary rank) and, prior to this paper, that had been a barrier to working with general algebras.

We identify *adjacency algebras of coherent configurations* as a promising family of algebras in the generalized framework. Coherent configurations are combinatorial objects that generalize groups and group actions; adjacency algebras are the analogue of group algebras and retain many of their important features. As with groups, coherent configurations support matrix multiplication when a natural combinatorial condition is satisfied, involving triangles of points in their underlying geometry.

Finally, we prove a closure property involving symmetric powers of adjacency algebras, which enables us to prove nontrivial bounds on ω using *commutative* coherent configurations and suggests that commutative coherent configurations may be sufficient to prove $\omega = 2$. Altogether, our results show that bounds on ω can be established by embedding large matrix multiplication instances into small commutative coherent configurations.

1 Introduction

Determining the exponent of matrix multiplication is one of the most fundamental unsolved problems in algebraic complexity. This quantity is the smallest number ω such that $n \times n$ matrix multiplication can be carried out using $n^{\omega+o(1)}$ arithmetic operations as n tends to infinity. Clearly $\omega \geq 2$, and it is widely believed that $\omega = 2$, but the best upper bound known is $\omega \leq 2.3727$ (due to Vassilevska Williams [20]). The importance of ω is by no means limited to matrix

multiplication, as ω also describes the asymptotic complexity of many other problems in linear algebra and graph theory (see Chapter 16 of [5]).

In the 43 years since Strassen’s original paper [18] gave the first improvement on the obvious exponent bound $\omega \leq 3$, there have been several major conceptual advances in the effort to obtain upper bounds on ω , each of which can informally be understood as relaxing the “rules of the game.” For example, Bini [2] showed that an upper bound on the *border rank* of a tensor implies an upper bound on its asymptotic rank. Indeed, there are useful examples of tensors with border rank strictly smaller than their rank, which led to improvements over Strassen’s original algorithm. Schönhage [16] showed how to convert upper bounds on the rank of the direct sum of several matrix multiplication tensors into an upper bound on ω , and his *asymptotic sum inequality* has played a crucial role in nearly all further advances. Strassen’s *laser method* [19] gave a way to convert non-matrix multiplication tensors (whose coarse structure contains a large diagonal, and whose components are all isomorphic to matrix multiplication tensors) into upper bounds on ω , and this method was used by Coppersmith and Winograd [9] as well as in the recent improvements of Davie and Stothers [17, 10] and Vassilevska Williams [20].

Here we introduce a further relaxation of the rules of the game, by studying a weighted version of matrix multiplication. Instead of computing the product AB of two matrices via

$$(AB)_{i,k} = \sum_j A_{i,j} B_{j,k},$$

we use

$$\sum_j \lambda_{i,j,k} A_{i,j} B_{j,k},$$

where the coefficients $\lambda_{i,j,k}$ are nonzero complex numbers. Of course, in certain cases weighted matrix multiplication is trivially equivalent to ordinary matrix multiplication. For example, if $\lambda_{i,j,k}$ can be written as $\alpha_{i,j} \beta_{j,k} \gamma_{k,i}$, then weighted matrix multiplication amounts to ordinary multiplication of matrices whose entries have been rescaled. However, rescaling does not yield an efficient equivalence for arbitrary weights.

We capture the complexity of weighted matrix multiplication via a new exponent ω_s , satisfying $2 \leq \omega_s \leq \omega$. It is the smallest real number for which there exist weights (depending on the dimensions of the matrices) such that weighted $n \times n$

*Microsoft Research New England, One Memorial Drive, Cambridge, MA 02142, cohn@microsoft.com

†Computing + Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, umans@cms.caltech.edu; research supported by NSF grants CCF-0846991 and CCF-1116111 and BSF grant 2010120.

matrix multiplication can be carried out in $n^{\omega_s+o(1)}$ arithmetic operations. The “s” stands for “support,” because we are dealing with tensors that have the same support as the matrix multiplication tensors.

In [8], we showed how to embed matrix multiplication into group algebra multiplication, and this methodology was used in [7] to prove strong bounds on ω . Replacing group algebras with more general algebras has always been an appealing generalization, and indeed the same approach works, except that it yields an embedding of *weighted* matrix multiplication. Thus, it gives an upper bound on ω_s , rather than ω . Prior to this paper, an upper bound on ω_s was of interest only by analogy with ω , and it was not known to imply anything about ω itself. Here, we overcome this obstacle by bounding ω in terms of ω_s , and we develop this embedding approach for a promising class of algebras.

Our main results are:

- (1) We prove that $\omega \leq (3\omega_s - 2)/2$. In particular, if $\omega_s \leq 2 + \varepsilon$, then $\omega \leq 2 + (3/2)\varepsilon$, so bounds for ω_s can be translated into bounds for ω with just a 50% penalty. Of course, that penalty is significant when ε is large, but our bound makes weighted matrix multiplication a viable approach for proving that $\omega = 2$ (as then $\varepsilon = 0$). This inequality between ω and ω_s can be proved using the laser method, but it does not seem to have been observed previously. We give a direct and self-contained proof in Section 3 as well as an explanation via the laser method. We also show that Boolean matrix multiplication has a randomized algebraic algorithm with running time $n^{\omega_s+o(1)}$, which avoids the 50% penalty.
- (2) We identify adjacency algebras of coherent configurations as a promising family of algebras. Coherent configurations are combinatorial objects that generalize groups and group actions; adjacency algebras are the analogue of group algebras and retain many of their important features. In particular, each adjacency algebra possesses a basis corresponding to an underlying geometry, and weighted matrix multiplication can be embedded when the coherent configuration satisfies a combinatorial condition involving triangles of points.
- (3) We prove a fundamental closure property of this class of algebras: any bound on ω_s obtained by applying the asymptotic sum inequality to independent embeddings of several weighted matrix multiplications can also be proved using a single embedding into a *symmetric power* of the algebra. Symmetric powers of adjacency algebras are themselves adjacency algebras, and this operation also preserves commutativity. Our results open the possibility of achieving $\omega = 2$ using commutative adjacency algebras, and we conjecture that commutative adjacency algebras suffice. In fact, that would follow from either of the two conjectures in [7].

A simple pigeonhole principle argument (Lemma 3.1 in [8]) shows that one cannot nontrivially embed a single matrix multiplication problem into a commutative group algebra. One might expect a similar barrier for commutative adjacency algebras, but the pigeonhole argument breaks down in this setting. Indeed, in this paper we prove nontrivial bounds on ω using commutative adjacency algebras in Theorem 5.6, by applying our machinery to the constructions from [7] (although we do not improve on the best known bounds). We should note that the simultaneous triple product property from [7] previously showed that one could avoid noncommutativity at the cost of having to deal with several independent embeddings. One could return to the setting of a single embedding using the wreath product construction (Theorem 7.1 in [7]), but this reintroduced noncommutativity, whereas working with coherent configurations rather than groups, as we do in this paper, avoids it completely.

The advantage of commutativity is that obtaining exponent bounds then amounts to a familiar type of task: embed as large an object as possible (here, a matrix multiplication instance) into as small an object as possible (here, a coherent configuration, with “size” measured by rank). By contrast, the noncommutative case involves a third quantity, namely the dimensions of the irreducible representations of the algebra.

2 Preliminaries and background

We define $[n] = \{1, 2, \dots, n\}$.

2.1 Tensors Our results will all be stated in terms of tensors. Recall that tensors are a generalization of vectors and matrices to higher orders. Tensor products of vector spaces form an elegant algebraic setting for the theory of tensors, but we will adopt the more concrete approach of representing tensors as multilinear forms. For example, the matrix with entries $A_{i,j}$ corresponds to the bilinear form $\sum_{i,j} A_{i,j} \hat{x}_i \hat{y}_j$, where \hat{x}_i and \hat{y}_j are formal variables, and we can represent a third-order tensor as $\sum_{i,j,k} A_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k$. We will use hats to make it clear which symbols denote formal variables. Applying invertible linear transformations to the sets of variables (here, $\{\hat{x}_i\}$, $\{\hat{y}_j\}$, and $\{\hat{z}_k\}$) yields an isomorphic tensor, but we cannot mix variables from different sets.

The *direct sum* $T \oplus T'$ of two tensors is simply their sum, if they have no variables in common (otherwise, first change variables to remove any overlap). For the *tensor product* $T \otimes T'$, if $T = \sum_{i,j,k} T_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k$ and $T' = \sum_{\ell,m,n} T'_{\ell,m,n} \hat{u}_\ell \hat{v}_m \hat{w}_n$, then

$$T \otimes T' = \sum_{i,j,k,\ell,m,n} T_{i,j,k} T'_{\ell,m,n} \hat{r}_{i,\ell} \hat{s}_{j,m} \hat{t}_{k,n},$$

with new variables $\hat{r}_{i,\ell}$, $\hat{s}_{j,m}$, and $\hat{t}_{k,n}$. In other words, we simply take the product of T and T' but combine the variables as illustrated above (e.g., $\hat{x}_i \hat{u}_\ell$ becomes $\hat{r}_{i,\ell}$). The direct sum and tensor product are defined only for tensors of the same order, and they preserve that order.

The rank $R(T)$ of a tensor T is one of its most important invariants. A nonzero tensor has rank 1 if it is the product of linear forms, and rank r if it is the sum of r rank 1 tensors but no fewer. In other words, $\sum_{i,j,k} T_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k$ has rank at most r if there are linear forms $\alpha_\ell(\hat{x})$, $\beta_\ell(\hat{y})$, and $\gamma_\ell(\hat{z})$ such that

$$\sum_{i,j,k} T_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k = \sum_{\ell=1}^r \alpha_\ell(\hat{x}) \beta_\ell(\hat{y}) \gamma_\ell(\hat{z}).$$

Tensor rank generalizes the concept of matrix rank, but it is more subtle. While matrices can be brought into a simple canonical form (row echelon form) in which their rank is visible, tensors cannot, because the symmetry group acting on them has far too low a dimension compared with the dimension of the space of tensors itself. Indeed, computing tensor rank is NP-hard [11].

2.2 Matrix multiplication in terms of tensors The *matrix multiplication tensor* $\langle \ell, m, n \rangle$ is the tensor

$$\sum_{i=1}^{\ell} \sum_{j=1}^m \sum_{k=1}^n \hat{x}_{i,j} \hat{y}_{j,k} \hat{z}_{k,i}.$$

Note that the coefficient of $\hat{z}_{k,i}$ singles out the $\hat{x}_{i,j} \hat{y}_{j,k}$ terms that occur in the (i, k) entry of the matrix product. It is easy to check that $\langle \ell, m, n \rangle \otimes \langle \ell', m', n' \rangle \cong \langle \ell \ell', m m', n n' \rangle$, which amounts to the assertion that block matrix multiplication computes the matrix product.

A low-rank expression for $\langle \ell, m, n \rangle$ specifies an efficient *bilinear algorithm* for computing the product of $\ell \times m$ and $m \times n$ matrices. In particular, it follows that $(\ell m n)^{\omega/3} \leq R(\langle \ell, m, n \rangle)$ (Proposition 15.5 in [5]).

In fact, although we have defined ω in terms of arbitrary algebraic algorithms, it is completely characterized by rank via

$$\omega = \inf\{\tau \in \mathbb{R} : R(\langle n, n, n \rangle) = O(n^\tau)\}$$

(Proposition 15.1 in [5]). In other words, bilinear algorithms have the same exponent as arbitrary algebraic algorithms. Thus, the entire subject of fast matrix multiplication can be reduced to bounding the rank of matrix multiplication tensors.

Schönhage’s *asymptotic sum inequality* [16] states that

$$(2.1) \quad (\ell_1 m_1 n_1)^{\omega/3} + \dots + (\ell_k m_k n_k)^{\omega/3} \leq R(\langle \ell_1, m_1, n_1 \rangle \oplus \dots \oplus \langle \ell_k, m_k, n_k \rangle),$$

and furthermore that the same holds for border rank (a relaxation of the notion of rank which will not play an important role in this paper). Thus, an unexpectedly efficient method for carrying out several independent matrix multiplications yields a bound on ω .

See [5] for further background on tensors, matrix multiplication, and algebraic complexity in general. It is important to keep in mind that the tensor manipulations all have implicit

algorithms behind them. In principle one could dispense with the tensor formalism completely, but it plays a valuable role in focusing attention on the central issues.

3 Matrix multiplication exponent bounds via s-rank

In this section, we show that an upper bound on what we call the “support rank”—or *s-rank*—of a matrix multiplication tensor implies an upper bound on ω . The *support* $\text{supp}(T)$ of a tensor T is the set of monomials that have nonzero coefficients. Of course this depends on the choice of basis and is therefore not an isomorphism invariant, and the same is true for concepts like s-rank that are defined in terms of it. However, basis dependence is not a difficulty in algebraic complexity. After all, any computational problem must specify a choice of basis for use in input and output, and writing a tensor as a multilinear form already involves an implicit choice of basis (the choice of variables).

DEFINITION 3.1. The *s-rank* $R_s(T)$ of a tensor T is the minimum rank of a tensor T' for which $\text{supp}(T) = \text{supp}(T')$.

Clearly s-rank can be no larger than rank. Here is a simple example that shows that s-rank can be dramatically smaller than both rank and border rank:

PROPOSITION 3.2. *The $n \times n$ matrix $J - I$, where J is the all ones matrix and I is the identity matrix, has rank n and border rank n , but s-rank equal to 2.*

Proof. Rank and border rank coincide for matrices (the matrices of rank at most r are characterized by determinantal conditions and thus form a closed set), and $J - I$ has rank n . However, consider the rank 1 matrix M defined by $M_{i,j} = \zeta^{i-j}$, where ζ is a primitive n -th root of unity. Then $M - J$ has the same support as $J - I$. Because M and J are both rank 1 matrices, the s-rank of $J - I$ is at most 2. It is also easy to see that no matrix with the same support as $J - I$ has rank 1, so the s-rank is exactly 2. \square

On the other hand, border rank can be smaller than s-rank, so these two relaxations of rank are incomparable:

PROPOSITION 3.3. *The tensor $T = \hat{x}_0 \hat{y}_0 \hat{z}_0 + \hat{x}_0 \hat{y}_1 \hat{z}_1 + \hat{x}_1 \hat{y}_0 \hat{z}_1$ has border rank 2 and s-rank 3.*

Proof. We refer to Bläser’s notes [4, p. 31] for the simple proof that the border rank is 2. To show that the s-rank is at least 3, we mimic his proof via the substitution method that the (ordinary) rank is 3. In a decomposition of any tensor T' with the same support as T into the sum of rank one tensors, one of the rank one tensors must depend on \hat{x}_1 . We can make this tensor zero by substituting a scalar multiple of \hat{x}_0 for \hat{x}_1 . After this substitution, T' still depends on \hat{y}_1 , so there is another rank one tensor in the decomposition that depends on \hat{y}_1 . We can make this tensor zero by substituting a scalar multiple of \hat{y}_0 for \hat{y}_1 . After both substitutions, T' still depends

on \hat{z}_0 , so there must be at least one more rank one tensor in the decomposition. The corresponding s-rank upper bound of 3 is trivial. \square

Like ordinary rank, s-rank is subadditive and submultiplicative in the sense that for tensors T and T' , we have $R_s(T \oplus T') \leq R_s(T) + R_s(T')$ and $R_s(T \otimes T') \leq R_s(T)R_s(T')$. For matrix multiplication tensors, we have $R_s(\langle \ell, m, n \rangle) = R_s(\langle \ell', m', n' \rangle)$ for every permutation (ℓ', m', n') of (ℓ, m, n) . In analogy to the exponent of matrix multiplication ω , we define ω_s , the s-rank exponent of matrix multiplication, as follows:

DEFINITION 3.4. The s-rank exponent of matrix multiplication, denoted ω_s , is defined by

$$\omega_s = \inf\{\tau \in \mathbb{R} : R_s(\langle n, n, n \rangle) = O(n^\tau)\}.$$

By comparison, the exponent ω can be defined in the same way, with ordinary rank replacing s-rank in the above expression. Since every tensor having the same support as $\langle n, n, n \rangle$ has n^2 linearly independent slices, $R_s(\langle n, n, n \rangle) \geq n^2$, and thus $2 \leq \omega_s \leq \omega$.

As one would expect, an s-rank upper bound implies an upper bound on ω_s . Here is the s-rank version of the standard proof:

PROPOSITION 3.5. For all ℓ, m, n , we have

$$(\ell mn)^{\omega_s/3} \leq R_s(\langle \ell, m, n \rangle).$$

Proof. Let $r = R_s(\langle \ell, m, n \rangle)$ and $M = \ell mn$. By symmetrizing, we have $R_s(\langle M, M, M \rangle) \leq r^3$, and then for all $N \geq 1$, by padding to the next largest power M^i of M ,

$$\begin{aligned} R_s(\langle N, N, N \rangle) &\leq R_s(\langle M^i, M^i, M^i \rangle) \\ &\leq r^{3i} \\ &= (M^i)^{3 \log_M r} \\ &= O(N^{3 \log_M r}). \end{aligned}$$

Thus, $\omega_s \leq 3 \log_M r$, from which the theorem follows. \square

We note that one can define the border s-rank of T to be the minimum border rank of tensors with the same support as T , and then by using Bini's argument [2], the above proposition holds with border s-rank in place of s-rank.

Whereas an upper bound on the rank of a matrix multiplication tensor implies a bilinear algorithm for matrix multiplication, an upper bound on the s-rank implies a bilinear algorithm for a *weighted* version of matrix multiplication: given matrices A and B , the algorithm computes values

$$C_{i,k} = \sum_j \lambda_{i,j,k} A_{i,j} B_{j,k},$$

where the weights $\lambda_{i,j,k}$ are certain nonzero scalars (depending on the construction used to attain a low s-rank). In other words, each entry of the result matrix C is a weighted inner product, with different weightings for the different inner products. There seems to be no obvious transformation to remove these weights.

As noted above, $2 \leq \omega_s \leq \omega$, so upper bounds on ω imply upper bounds on ω_s (and $\omega = 2$ implies $\omega_s = 2$). Theorem 3.6 below shows that upper bounds on ω_s imply upper bounds on ω , and indeed $\omega_s = 2$ implies $\omega = 2$. Thus s-rank is a useful relaxation of rank, when trying to bound the exponent of matrix multiplication.

THEOREM 3.6. The exponents ω and ω_s satisfy

$$\omega \leq (3\omega_s - 2)/2.$$

In other words, $\omega_s \leq 2 + \varepsilon$ implies $\omega \leq 2 + (3/2)\varepsilon$.

Proof. By the definition of ω_s , we have $R_s(\langle n, n, n \rangle) = n^{\omega_s + o(1)}$. For a given value of n , let T be the trilinear form corresponding to this (weighted) $n \times n$ matrix multiplication:

$$T = \sum_{a,b,c \in [n]} \lambda_{a,b,c} \hat{x}_a \hat{y}_b \hat{z}_{c,a}$$

with $0 \neq \lambda_{a,b,c} \in \mathbb{C}$ for all a, b, c . Let

$$S \subseteq \Delta_n = \{(s_1, s_2, s_3) : s_1, s_2, s_3 \in [n] \text{ and } s_1 + s_2 + s_3 = n + 2\}$$

be a *triangle-free set*, as defined in Section 6.2 of [7]. Such a set has the property that if $s, t, u \in S$ satisfy $s_1 = t_1, t_2 = u_2$, and $u_3 = s_3$ then $s = t = u$. In [7] we gave a simple construction of triangle-free sets S with $|S| = n^{2-o(1)}$. Let T' be the trilinear form corresponding to $|S|$ independent $n^2 \times n^2$ matrix multiplications; i.e.,

$$T' = \sum_{s \in S, i,j,k \in [n]^2} \hat{u}_{s,i,j} \hat{v}_{s,j,k} \hat{w}_{s,k,i}.$$

We will show that T' is a restriction of the tensor power $T^{\otimes 3}$, which is given by

$$T^{\otimes 3} = \sum_{a,b,c \in [n]^3} \lambda_{a_1,b_1,c_1} \lambda_{a_2,b_2,c_2} \lambda_{a_3,b_3,c_3} \hat{x}_a \hat{y}_b \hat{z}_{c,a}.$$

In other words, we will show that T' can be obtained by substituting variables in $T^{\otimes 3}$, which implies that $R(T') \leq R(T^{\otimes 3})$. To do so, define (for $s, t, u \in S$ and $i, i', j, j', k, k' \in [n]^2$)

$$\begin{aligned} \hat{u}_{s,i,j'} &= \lambda_{i_2,j'_1,s_2} \hat{x}_{(i_1,i_2,s_3),(s_1,j'_1,j'_2)} \\ \hat{v}_{t,j,k'} &= \lambda_{t_3,j_2,k'_2} \hat{y}_{(t_1,j_1,j_2),(k'_1,t_2,k'_2)} \\ \hat{w}_{u,k,i'} &= \lambda_{i'_1,u_1,k_1} \hat{z}_{(k_1,u_2,k_2),(i'_1,i'_2,u_3)} \end{aligned}$$

and set the $\hat{x}, \hat{y}, \hat{z}$ variables not mentioned in these equations equal to zero. Under this change of variables, we will see

that $T^{\otimes 3}$ becomes exactly the tensor T' . To check this, we must verify that upon substituting the $\hat{u}, \hat{v}, \hat{w}$ variables for the $\hat{x}, \hat{y}, \hat{z}$ variables in $T^{\otimes 3}$ according to the above formulas, the coefficient of $\hat{u}_{s,i,j'} \hat{v}_{t,j,k'} \hat{w}_{u,k,i'}$ is 1 if $s = t = u$, $i = i'$, $j = j'$, and $k = k'$, and it is 0 otherwise. Since the support of $T^{\otimes 3}$ is the same as the support of $\langle n^3, n^3, n^3 \rangle$, the monomial

$$\hat{x}_{(i_1, i_2, s_3), (s_1, j'_1, j'_2)} \hat{y}_{(t_1, j_1, j_2), (k'_1, t_2, k'_2)} \hat{z}_{(k_1, u_2, k_2), (i'_1, i'_2, u_3)}$$

in $T^{\otimes 3}$ has a nonzero coefficient if and only if

$$\begin{aligned} (s_1, j'_1, j'_2) &= (t_1, j_1, j_2) \\ (k'_1, t_2, k'_2) &= (k_1, u_2, k_2) \\ (i'_1, i'_2, u_3) &= (i_1, i_2, s_3). \end{aligned}$$

This happens if and only if $i = i'$, $j = j'$, $k = k'$, $s_1 = t_1$, $t_2 = u_2$, and $u_3 = s_3$, and by the definition of a triangle-free set the last three conditions imply $s = t = u$. The coefficient in $T^{\otimes 3}$ of

$$\hat{x}_{(i_1, i_2, s_3), (s_1, j_1, j_2)} \hat{y}_{(s_1, j_1, j_2), (k_1, s_2, k_2)} \hat{z}_{(k_1, s_2, k_2), (i_1, i_2, s_3)}$$

is $\lambda_{i_1, s_1, k_1} \lambda_{i_2, j_1, s_2} \lambda_{s_3, j_2, k_2}$, which exactly corresponds to the $\lambda_{i_2, j'_1, s_2} \lambda_{t_3, j_2, k'_2} \lambda_{i'_1, u_1, k_1}$ factor from the definition of $\hat{u}, \hat{v}, \hat{w}$, so the coefficient of $\hat{u}_{s,i,j'} \hat{v}_{t,j,k'} \hat{w}_{u,k,i'}$ after the substitution is 1.

Thus the ordinary rank of the direct sum of $|S| = n^{2-o(1)}$ independent $n^2 \times n^2$ matrix multiplications is at most $(n^{\omega_s+o(1)})^3$, and applying the asymptotic sum inequality (2.1), we get

$$n^{2-o(1)} n^{2\omega} \leq (n^{\omega_s+o(1)})^3.$$

Taking logarithms and letting n go to infinity, we get $2 + 2\omega \leq 3\omega_s$, as desired. \square

Theorem 3.6 can also be proved using the laser method, in particular using Proposition 7.3 from [19] (Proposition 15.32 in [5]), although this consequences does not seem to have been observed in the literature. Here is a sketch of the proof. By definition, there exist tensors having the same support as $\langle n, n, n \rangle$, with rank $n^{\omega_s+o(1)}$. Adopting the language of [5], such a tensor has a direct sum decomposition D whose D -support is isomorphic to $\langle 1, n, 1 \rangle$ and whose D -components are each *isomorphic* to $\langle n, 1, n \rangle$ (it is a special feature of an outer-product tensor, which has only a single 1 in each slice, that all tensors with the same support are isomorphic). Proposition 15.32 in [5] then implies that $n^2 n^{2\omega} \leq (n^{\omega_s+o(1)})^3$, which yields the same bound as Theorem 3.6.

It is an interesting open problem to improve the conclusion of Theorem 3.6 to $\omega \leq \omega_s$. In the preceding paragraph, the D -component isomorphism being used is very special (it corresponds to a *diagonal* change of basis), so one might hope to avoid the loss coming from machinery that handles arbitrary isomorphisms.

In fact, this loss *can* be avoided entirely when one is interested in the simpler problem of Boolean matrix multiplication, i.e., matrix multiplication over the Boolean semiring with “and” as multiplication and “or” as addition. The next theorem describes how to use weighted multiplication directly to obtain an algebraic algorithm for Boolean matrix multiplication.

Given $n \times n$ Boolean matrices A and B , let A' be the obvious lift of A to a 0/1 complex matrix, and define B' the same way but with a random choice of 1 or 2 for each nonzero entry.

THEOREM 3.7. *There is an algebraic algorithm running in $n^{\omega_s+o(1)}$ operations that computes from A', B' an $n \times n$ matrix C with the following property: the (i, j) entry of C is 0 if the (i, j) entry of the Boolean matrix product of A and B is 0; otherwise it is nonzero with probability at least $1/2$.*

The procedure may be repeated $O(\log n)$ times to obtain all entries of the Boolean product of A and B with high probability.

Proof. Using a bilinear algorithm, one can compute a suitably weighted product of A' and B' in $n^{\omega_s+o(1)}$ operations. We get a result matrix whose (i, j) entry is

$$\sum_{\ell} \lambda_{i,\ell,j} A'_{i,\ell} B'_{\ell,j},$$

where $\lambda_{i,\ell,j} \neq 0$. When the (i, j) entry of the Boolean matrix product of A and B is zero, this value is clearly also zero; otherwise, it equals $\sum_{\ell \in L} \lambda_{i,\ell,j} r_{\ell}$ for a nonempty set L and each r_{ℓ} chosen randomly from $\{1, 2\}$. For a given $\ell \in L$, there is a unique value of r_{ℓ} making this sum zero, so the probability that it vanishes is at most $1/2$. \square

If the weights arising in the above proof are all positive (as they are for all the s -rank bounds we derive in this paper), then no randomness is needed, as A' and B' can both be taken to be the obvious lifts of A and B to 0/1 matrices.

Finally, we note that all of the manipulations used in the matrix multiplication literature for converting bounds on the rank of certain “basic” tensors into bounds on the rank of the matrix multiplication tensor also work with s -rank in place of rank. So, for example, s -rank bounds on the partial matrix multiplication tensor of Bini, Capovani, Lotti, and Romani [3], the basic tensor used by Schönhage [16], the basic tensor in Strassen’s laser method paper [19], or any of the basic tensors introduced by Coppersmith and Winograd [9] eventually yield an s -rank bound on a matrix multiplication tensor by simply following the known proofs. However, for most of these basic tensors with explicit tensor decompositions, there are matching lower bounds on the rank via the *substitution method* (e.g., the proof of Proposition 3.3). Substitution method lower bounds also prove lower bounds

on s-rank, so there does not seem to be an opportunity for an easy improvement from switching to s-rank. However, an improvement by switching to border s-rank might be possible. As a concrete example, we do not know whether the border s-rank, or even just the s-rank, of $\langle 2, 2, 2 \rangle$ is 6 or 7.

4 Matrix multiplication via coherent configurations

In this section, we describe how to embed matrix multiplication into algebra multiplication. To do so, it is helpful to have a basis with considerable combinatorial structure. As mentioned in the introduction, adjacency algebras of coherent configurations are a promising family of algebras to use here. In Subsections 4.2 and 4.3, we review the basic theory of coherent configurations and their adjacency algebras, and in Subsection 4.4 we specialize our general theory to this setting.

4.1 Realizing matrix multiplication in algebras We can bound the s-rank of matrix multiplication by restricting the structural tensor of an algebra. Let A be a finite-dimensional complex algebra, and let u_1, \dots, u_r be a basis for A . Then there are coefficients $\lambda_{i,j,k}$ such that

$$u_i u_j = \sum_{k=1}^r \lambda_{i,j,k} u_k;$$

they are called the *structure constants* of A with respect to this basis. The *structural tensor* is the trilinear form

$$\sum_{i,j,k} \lambda_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k.$$

It is isomorphic to the multiplication tensor (i.e., the element of $A^* \otimes A^* \otimes A$ corresponding to the multiplication map from $A \otimes A$ to A). More generally, if we use any three bases u_1, \dots, u_r , v_1, \dots, v_r , and w_1, \dots, w_r for A and define the coefficients by

$$u_i v_j = \sum_{k=1}^r \lambda_{i,j,k} w_k,$$

then the corresponding tensor is isomorphic to the structural tensor.

DEFINITION 4.1. Let A be an r -dimensional complex algebra with structure constants $\lambda_{i,j,k}$ corresponding to some choice of bases. We say A *realizes* $\langle \ell, m, n \rangle$ if there exist three injective functions

$$\alpha: [\ell] \times [m] \rightarrow [r], \quad \beta: [m] \times [n] \rightarrow [r], \quad \gamma: [n] \times [\ell] \rightarrow [r]$$

such that

$$\lambda_{\alpha(a,b'), \beta(b,c'), \gamma(c,a')} \neq 0$$

if and only if $a = a'$, $b = b'$, and $c = c'$.

Note that this definition depends on the choice of basis. We will typically suppress the choice of basis in the notation, because the algebras we deal with later in the paper will always come with a standard basis.

One might reasonably use the term “s-realize” instead of “realize” in Definition 4.1. We have chosen to use the simpler term, rather than reserving it for strict realization involving only structure constants that are 0 or 1, because we know of few interesting examples of strict realization beyond group algebras (where the notions coincide).

PROPOSITION 4.2. *If an algebra A realizes $\langle \ell, m, n \rangle$, then the s-rank of $\langle \ell, m, n \rangle$ is at most the rank of the structural tensor for A .*

Proof. Suppose A realizes $\langle \ell, m, n \rangle$ via α, β, γ , and consider the structural tensor

$$\sum_{i,j,k} \lambda_{i,j,k} \hat{x}_i \hat{y}_j \hat{z}_k.$$

Define $\hat{u}_{a,b'} = \hat{x}_{\alpha(a,b')}$, $\hat{v}_{b,c'} = \hat{y}_{\beta(b,c')}$, and $\hat{w}_{c,a'} = \hat{z}_{\gamma(c,a')}$; furthermore, set $\hat{x}_i = 0$ when i is not in the image of α , $\hat{y}_j = 0$ when j is not in the image of β , and $\hat{z}_k = 0$ when k is not in the image of γ . Under this change of variables, the structural tensor becomes

$$\sum_{a,a',b,b',c,c'} \lambda_{\alpha(a,b'), \beta(b,c'), \gamma(c,a')} \hat{u}_{a,b'} \hat{v}_{b,c'} \hat{w}_{c,a'},$$

and by assumption the terms vanish unless $a = a'$, $b = b'$, and $c = c'$. Thus,

$$\sum_{a,b,c} \lambda_{\alpha(a,b), \beta(b,c), \gamma(c,a)} \hat{u}_{a,b} \hat{v}_{b,c} \hat{w}_{c,a}$$

has rank at most that of the structural tensor. This new tensor is a weighting of the matrix multiplication tensor $\langle \ell, m, n \rangle$, so the s-rank of $\langle \ell, m, n \rangle$ is at most the rank of the structural tensor. \square

Note that this proof in fact gives a very simple algorithm for reducing a weighted matrix multiplication to an algebra multiplication, along the lines of the reduction in [8].

Recall that an algebra is semisimple if it is a product of matrix algebras. In other words, it is semisimple if there are *character degrees* d_1, \dots, d_t so that

$$A \cong \mathbb{C}^{d_1 \times d_1} \times \dots \times \mathbb{C}^{d_t \times d_t}.$$

In that case, the structural tensor is isomorphic to $\langle d_1, d_1, d_1 \rangle \oplus \dots \oplus \langle d_t, d_t, d_t \rangle$.

PROPOSITION 4.3. *If a semisimple algebra A with character degrees d_1, \dots, d_t realizes $\langle \ell, m, n \rangle$, then*

$$(\ell mn)^{\omega_s/3} \leq d_1^\omega + \dots + d_t^\omega.$$

This proposition also involves a natural algorithm, which reduces a weighted matrix multiplication to a collection of unweighted matrix multiplications.

Proof. For each $\varepsilon > 0$, there is a constant C such that $R(\langle d, d, d \rangle) \leq Cd^{\omega+\varepsilon}$ for all d . It follows that

$$(\ell mn)^{\omega_s/3} \leq R_s(\langle \ell, m, n \rangle) \leq Cd_1^{\omega+\varepsilon} + \dots + Cd_t^{\omega+\varepsilon},$$

but the C and ε are problematic. To remove them, we will use the trick of computing the asymptotic rank for high tensor powers. The algebra $A^{\otimes N}$ realizes a weighted version of $\langle \ell, m, n \rangle^{\otimes N} = \langle \ell^N, m^N, n^N \rangle$, and it has character degrees given by N -fold products $d_{i_1} \dots d_{i_N}$. Thus,

$$(\ell mn)^{N\omega_s/3} \leq C(d_1^{\omega+\varepsilon} + \dots + d_t^{\omega+\varepsilon})^N.$$

Now taking N -th roots and letting N tend to infinity yields

$$(\ell mn)^{\omega_s/3} \leq d_1^{\omega+\varepsilon} + \dots + d_t^{\omega+\varepsilon},$$

and because this holds for all $\varepsilon > 0$ it also holds for $\varepsilon = 0$ by continuity. \square

Definition 4.1 generalizes the triple product property from [8]. Recall that three subsets S, T, U of a group satisfy the *triple product property* if

$$\begin{aligned} s^{-1}s't^{-1}t'u^{-1}u' &= 1 \\ \Leftrightarrow \\ s = s', t = t', u = u' \end{aligned}$$

holds for $s, s' \in S$, $t, t' \in T$, and $u, u' \in U$. To see why Definition 4.1 is a generalization, suppose A is the group algebra of a finite group, and choose the group elements themselves as a basis. Then $\alpha(a, b')$, $\beta(b, c')$, and $\gamma(c, a')$ correspond to group elements $g_{a,b'}$, $h_{b,c'}$ and $k_{a',c}$ such that

$$(4.1) \quad g_{a,b'}h_{b,c'} = k_{a',c}$$

if and only if $a = a'$, $b = b'$, and $c = c'$. We wish to find group elements s_a, t_b, u_c such that $g_{a,b} = s_a t_b^{-1}$, $h_{b,c} = t_b u_c^{-1}$, and $k_{a,c} = s_a u_c^{-1}$; then $\{s_a\}, \{t_b\}, \{u_c\}$ satisfy the triple product property. To find these group elements, fix b_0 , and let $s_a = g_{a,b_0}$ and $u_c = h_{b_0,c}^{-1}$. Then $s_a u_c^{-1} = k_{a,c}$ automatically. Furthermore, (4.1) implies that $g_{a,b}^{-1} g_{a,b_0} = h_{b_0,c} h_{b,c}^{-1}$, with this group element being independent of a and c . Calling it t_b completes the construction.

4.2 Coherent configurations Coherent configurations are remarkable structures that unify much of group theory and algebraic combinatorics [12, 13, 14]. A *coherent configuration* of rank r is a finite set \mathcal{C} , whose elements are called *points*, with a partition of \mathcal{C}^2 into subsets R_1, R_2, \dots, R_r called *classes* such that

(1) the diagonal $\{(x, x) : x \in \mathcal{C}\}$ is the union of some of the classes,

(2) for each $i \in [r]$ there exists $i^* \in [r]$ such that $R_i^* = R_{i^*}$, where

$$R_i^* = \{(b, a) : (a, b) \in R_i\},$$

and

(3) there exist integers $p_{i,j}^k$ for $i, j, k \in [r]$ such that for all $x, y \in \mathcal{C}$ with $(x, y) \in R_k$,

$$\#\{z \in \mathcal{C} : (x, z) \in R_i \text{ and } (z, y) \in R_j\} = p_{i,j}^k.$$

We say \mathcal{C} is *symmetric* if $R_i^* = R_i$ for all i and *commutative* if $p_{i,j}^k = p_{j,i}^k$ for all i, j, k . (Symmetry implies commutativity, but not vice versa.) The numbers $p_{i,j}^k$ are called the *intersection numbers* of the configuration. The configuration is an *association scheme* if the diagonal is itself one of the classes. It is easily proved that a commutative coherent configuration must be an association scheme [14, p. 14].

Every finite group G defines an association scheme, with G as its set of points and G^2 partitioned into subsets $R_g = \{(h, hg) : h \in G\}$ with $g \in G$. Then for $g, h, k \in G$,

$$p_{g,h}^k = \begin{cases} 1 & \text{if } gh = k, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

The intersection numbers encode the multiplication table of the group, so the group and the corresponding association scheme are fully equivalent structures. Note that this association scheme is commutative iff G is, while it is symmetric iff $g = g^{-1}$ for all $g \in G$. One can show that an association scheme comes from a group in this way if and only if all its intersection numbers are at most 1.

More generally, suppose G acts on a finite set X . Then partitioning X^2 into the orbits of G under the diagonal action defines a coherent configuration, called a *Schurian coherent configuration*. It is an association scheme iff G acts transitively on X .

Many important examples in combinatorics fit into this framework. For example, the Hamming scheme consists of the points in $\{0, 1\}^n$ with classes defined by Hamming distance. From a group-theoretic perspective, it is the Schurian association scheme defined by the action of the semidirect product $S_n \ltimes (\mathbb{Z}/2\mathbb{Z})^n$ on $\{0, 1\}^n$, although this formulation is excessive for most purposes.

If G acts transitively on X , then we can identify X with G/H , where H is the stabilizer of a point in X . Note that G/H is not a group unless H is a normal subgroup, but it is always an association scheme. In certain cases, called *Gelfand pairs* (G, H) , the quotient G/H is a commutative association scheme (although the groups G and H will typically not be commutative). For example, this occurs for the Hamming scheme.

There are also numerous combinatorial examples of association schemes and coherent configurations that do not come from symmetry groups. For example, strongly regular graphs are the same thing as symmetric association schemes of rank 3. More generally, every distance-regular graph is an association scheme when the classes are defined by the graph metric. Some of these graphs are Schurian association schemes, but many are not.

A *fusion* of a coherent configuration \mathcal{C} is a configuration \mathcal{C}' with the same set of points and with the classes of \mathcal{C}' all given by unions of classes of \mathcal{C} . (Note that this must be done carefully, since taking arbitrary unions will generally not yield a coherent configuration.) Another important construction is the *direct product*: given two coherent configurations \mathcal{C} and \mathcal{C}' , their product $\mathcal{C} \times \mathcal{C}'$ has the direct product of their point sets as its point set, with the class of (c_1, c'_1) and (c_2, c'_2) determined by the class of (c_1, c_2) in \mathcal{C} and that of (c'_1, c'_2) in \mathcal{C}' . The *symmetric power* $\text{Sym}^k \mathcal{C}$ is the fusion scheme formed by fusing the classes of the direct power \mathcal{C}^k under the action of the symmetric group S_k on the factors.

4.3 The adjacency algebra Every coherent configuration has an associated algebra, which plays the same role as the group algebra of a group. Let A_1, \dots, A_r be the adjacency matrices of the relations R_1, \dots, R_r . In other words, A_i is indexed by \mathcal{C} , with

$$(A_i)_{x,y} = \begin{cases} 1 & \text{if } (x,y) \in R_i, \text{ and} \\ 0 & \text{otherwise} \end{cases}$$

for $x, y \in \mathcal{C}$. The *adjacency algebra* $\mathbb{C}[\mathcal{C}]$ of \mathcal{C} is the complex algebra generated by these adjacency matrices. (Note that it contains the identity because the diagonal is a union of classes.) An easy calculation shows that

$$A_i A_j = \sum_k p_{i,j}^k A_k,$$

so $\mathbb{C}[\mathcal{C}]$ is spanned by A_1, \dots, A_r . It is a commutative algebra if and only if \mathcal{C} is commutative.

The adjacency algebra is closed under the conjugate transpose, so it is a semisimple algebra (see, for example, Theorem 3.2 in [6]). Thus, there exist *character degrees* d_1, \dots, d_k such that

$$\mathbb{C}[\mathcal{C}] \cong \mathbb{C}^{d_1 \times d_1} \times \dots \times \mathbb{C}^{d_k \times d_k}.$$

Of course, $d_1^2 + \dots + d_k^2$ must equal the dimension of $\mathbb{C}[\mathcal{C}]$, which is the rank of \mathcal{C} . The adjacency algebra of a commutative coherent configuration of rank r is isomorphic to \mathbb{C}^r .

The structural tensor of $\mathbb{C}[\mathcal{C}]$ is

$$\sum_{i,j,k \in [r]} p_{i,j}^k \hat{x}_i \hat{y}_j \hat{z}_k,$$

but (as we will see shortly) it is often convenient to use

$$\sum_{i,j,k \in [r]} p_{i,j}^{k*} \hat{x}_i \hat{y}_j \hat{z}_k$$

instead. This isomorphic tensor simply amounts to reordering the variables \hat{z}_k . Note that the rank r of the coherent configuration is not necessarily the same as the rank of the structural tensor: they are equal if and only if the configuration is commutative.

4.4 Embedding matrix multiplication into an adjacency algebra Let \mathcal{C} be a coherent configuration of rank r , with notation as in the previous subsection.

DEFINITION 4.4. Three classes i, j, k form a *triangle* if there exist points x, y, z such that $(x, y) \in R_i$, $(y, z) \in R_j$, and $(z, x) \in R_k$.

In terms of intersection numbers, classes i, j, k form a triangle iff $p_{i,j}^{k*} > 0$. (Note that we use k^* instead of k to switch the order of x and z .) This is why we prefer to use k^* instead of k in the structural tensor: otherwise, the cyclic symmetry among x, y, z is broken.

DEFINITION 4.5. A coherent configuration \mathcal{C} of rank r *realizes* $\langle \ell, m, n \rangle$ if there exist three injective functions

$$\alpha: [\ell] \times [m] \rightarrow [r], \quad \beta: [m] \times [n] \rightarrow [r], \quad \gamma: [n] \times [\ell] \rightarrow [r]$$

such that $\alpha(a, b'), \beta(b, c'), \gamma(c, a')$ form a triangle iff $a = a'$, $b = b'$, and $c = c'$.

Of course this definition assumes a fixed numbering of the classes in \mathcal{C} . It amounts to the general definition of realization in an algebra, specialized to our choice of structural tensor.

EXAMPLE 4.6. As a simple example, let \mathcal{C} be the coherent configuration on n points for which every pair of points defines a distinct class. If we index the classes with pairs of points, then (a, b') , (b, c') , and (c, a') form a triangle if and only if $a = a'$, $b = b'$, and $c = c'$, so $\mathbb{C}[\mathcal{C}]$ trivially realizes $\langle n, n, n \rangle$. As one might expect from such a trivial example, the embedding yields no benefit for matrix multiplication, because in fact $\mathbb{C}[\mathcal{C}] \cong \mathbb{C}^{n \times n}$.

In the next section we will construct less trivial examples. In the meantime, we note the following proposition, which works out the conditions for a coherent configuration arising from a group action to realize matrix multiplication.

PROPOSITION 4.7. *Let G be a finite group acting on a set X , and let \mathcal{C} be the corresponding Schurian coherent configuration. Suppose there exist subsets $A, B, C \subseteq X$ such that for all $f, g, h \in G$ and all $a \in A, b \in B$, and $c \in C$,*

$$\begin{aligned} & \text{if } fa \in A, gb \in B, hc \in C \text{ and } fgh = 1, \\ & \text{then } fa = a, gb = b, \text{ and } hc = c. \end{aligned}$$

Then \mathcal{C} realizes $\langle |A|, |B|, |C| \rangle$.

Proof. Recall that the classes of \mathcal{C} are the orbits of G on X^2 . If we identify A with $[|A|]$, etc., then we realize $\langle |A|, |B|, |C| \rangle$ via maps α, β, γ such that for $a \in A$ and $b' \in B$, $\alpha(a, b')$ is the orbit of (a, b') in X^2 , etc. The map α is injective, because $\alpha(a, b) = \alpha(a', b')$ implies $fa = a'$ and $fb = b'$ for some $f \in G$, and the hypothesis then implies $a = a'$ and $b = b'$ (take $g = f^{-1}$ and $h = 1$); the maps β and γ are injective by the same argument. Now we wish to show that the classes of (a, b') , (b, c') , and (c, a') form a triangle if and only if $a = a'$, $b = b'$, and $c = c'$ (where $a, a' \in A$, $b, b' \in B$, and $c, c' \in C$). Saying they form a triangle means there exist $x, y, z \in X$ and $s, t, u \in G$ such that $(x, y) = (sa, sb')$, $(y, z) = (tb, tc')$, and $(z, x) = (uc, ua')$. If we set $f = u^{-1}s$, $g = s^{-1}t$, and $h = t^{-1}u$, then $fgh = 1$ and $fa = a'$, $gb = b'$, and $hc = c'$. Now by hypothesis we have $a = a'$, $b = b'$, and $c = c'$, as desired. \square

If we let G act on itself by left translation, then the hypothesis of Proposition 4.7 simply asserts that A , B , and C satisfy the triple product property from [8]. Thus, the proposition gives a natural generalization of the triple product property from groups to group actions.

5 Simultaneous embeddings and symmetric powers

Our best construction techniques so far are all based on realizing several independent matrix multiplications simultaneously; this was called the *simultaneous triple product property* in [7]. In a coherent configuration, the definition amounts to the following (and of course one can give an analogous definition in any algebra):

DEFINITION 5.1. A coherent configuration \mathcal{C} of rank r realizes $\oplus_i \langle \ell_i, m_i, n_i \rangle$ if there exist injective functions

$$\begin{aligned}\alpha_i &: [\ell_i] \times [m_i] \rightarrow [r] \\ \beta_i &: [m_i] \times [n_i] \rightarrow [r] \\ \gamma_i &: [n_i] \times [\ell_i] \rightarrow [r]\end{aligned}$$

such that $\alpha_i(a, b')$, $\beta_j(b, c')$, $\gamma_k(c, a')$ form a triangle iff $i = j = k$ and $a = a'$, $b = b'$, and $c = c'$.

If \mathcal{C} realizes $\langle \ell_1, m_1, n_1 \rangle \oplus \cdots \oplus \langle \ell_k, m_k, n_k \rangle$ and T is its structural tensor, then

$$R_s(\langle \ell_1, m_1, n_1 \rangle \oplus \cdots \oplus \langle \ell_k, m_k, n_k \rangle) \leq R(T).$$

One can imitate the proof of the asymptotic sum inequality to show that

$$(5.1) \quad (\ell_1 m_1 n_1)^{\omega_s/3} + \cdots + (\ell_k m_k n_k)^{\omega_s/3} \leq R_s(\langle \ell_1, m_1, n_1 \rangle \oplus \cdots \oplus \langle \ell_k, m_k, n_k \rangle),$$

from which one can deduce bounds on ω_s . Instead, in this section we will develop an efficient algebraic method to combine these independent matrix multiplication realizations

into one. It will yield the same bound, but also show that this bound is achieved by realizing a single matrix multiplication in a coherent configuration. First, we give an example. This example is extremal, because it realizes the direct sum of $n^{1-o(1)}$ copies of $\langle n, n, n \rangle$ via a coherent configuration of rank n^3 , and this cannot be done with rank less than $n^{3-o(1)}$ (because the images of the embeddings must be disjoint). If the coherent configuration were commutative, then we could conclude that $\omega = 2$, but it is far from commutative.

EXAMPLE 5.2. Let \mathcal{C} be the coherent configuration corresponding to the diagonal action of $\mathbb{Z}/n\mathbb{Z}$ on $(\mathbb{Z}/n\mathbb{Z})^2$, and let $S \subseteq \mathbb{Z}/n\mathbb{Z}$ be a set of size $|S| = n^{1-o(1)}$ containing no three-term arithmetic progression [15]. We can index the classes in \mathcal{C} as

$$R_{(a,b,c)} = \{((s, s+a), (s+a+b, s+a+b+c)) : s \in \mathbb{Z}/n\mathbb{Z}\},$$

with $a, b, c \in \mathbb{Z}/n\mathbb{Z}$. Then \mathcal{C} realizes $\oplus_{i \in S} \langle n, n, n \rangle$ via maps $\alpha_i, \beta_i, \gamma_i$ defined for $i \in S$ by

$$\begin{aligned}\alpha_i(x, y) &= (x, i-x, y) \\ \beta_i(y, z) &= (y, i-y, z) \\ \gamma_i(z, x) &= (z, -2i-z, x).\end{aligned}$$

Specifically, it is not hard to check that $(x, i-x, y')$, $(y, j-y, z')$, $(z, -2k-z, x')$ form a triangle if and only if $x = x'$, $y = y'$, $z = z'$, and $i+j = 2k$ (in which case $i = j = k$ because S contains no three-term arithmetic progressions). However, this example does not prove any nontrivial bound on ω , because in fact the character degrees of \mathcal{C} are all equal to n (repeated n times).

As promised, we now give a constructive proof of (5.1). The proof converts a coherent configuration that realizes several independent matrix multiplications into a single coherent configuration that realizes a *single* matrix multiplication. Moreover, the resulting coherent configuration is commutative if the original one was. Because the proof actually constructs a coherent configuration rather than just deducing the bound on ω_s , we can use it to obtain *commutative* coherent configurations that prove nontrivial bounds on ω . This establishes one of the main points of the paper: that the noncommutativity that was necessary in the group-theoretic approach can be avoided in the generalization to coherent configurations.

We also find that a consequence of either of the two main conjectures of [7] is that *commutative* coherent configurations suffice to prove $\omega = 2$. This raises our hope that one could find commutative coherent configurations of rank $n^{2+o(1)}$ that realize $\langle n, n, n \rangle$ and thus prove $\omega = 2$.

The main idea of the proof is to take symmetric powers, as described next:

THEOREM 5.3. *Let \mathcal{C} be a coherent configuration of rank r that realizes $\bigoplus_{i=1}^k \langle \ell_i, m_i, n_i \rangle$. Then the symmetric power $\text{Sym}^k \mathcal{C}$ realizes $\langle \prod_i \ell_i, \prod_i m_i, \prod_i n_i \rangle$ and has rank $\binom{r+k-1}{k}$.*

Proof. The classes R_I of the k -fold direct product of \mathcal{C} are indexed by vectors $I \in [r]^k$. The symmetric group S_k acts on $[r]^k$ by permuting the k coordinates, and the orbits of this action naturally correspond to the k -multisubsets of $[r]$. Recall that $\text{Sym}^k \mathcal{C}$ is the fusion configuration with a class for each orbit; i.e., for each k -multisubset S of r , we have a class R'_S that is the union of R_I over all I in the orbit corresponding to S . The rank of $\text{Sym}^k \mathcal{C}$ is the number of distinct orbits (equivalently, the number of distinct k -multisubsets of $[r]$), which is $\binom{r+k-1}{k}$.

Set $L = \prod_i \ell_i$, $M = \prod_i m_i$, and $N = \prod_i n_i$. Now, \mathcal{C}^k realizes $\langle L, M, N \rangle$, so there exist injective functions

$$\begin{aligned} \alpha: [L] \times [M] &\rightarrow [r]^k \\ \beta: [M] \times [N] &\rightarrow [r]^k \\ \gamma: [N] \times [L] &\rightarrow [r]^k \end{aligned}$$

satisfying the conditions of Definition 4.1; specifically,

$$\begin{aligned} \alpha(A, B) &= (\alpha_1(A_1, B_1), \dots, \alpha_k(A_k, B_k)) \\ \beta(B, C) &= (\beta_1(B_1, C_1), \dots, \beta_k(B_k, C_k)) \\ \gamma(C, A) &= (\gamma_1(C_1, A_1), \dots, \gamma_k(C_k, A_k)), \end{aligned}$$

where \mathcal{C} realizes $\bigoplus_{i=1}^k \langle \ell_i, m_i, n_i \rangle$ via $\alpha_i, \beta_i, \gamma_i$.

We claim that in fact α, β, γ are injective even in the fusion configuration, where we collapse the orbits. For suppose that $\alpha(A, B) = \pi \alpha(A', B')$ for some $\pi \in S_k$. Then π must be the identity since the maps α_i have disjoint images in $[r]$ (this follows immediately from Definition 5.1), and then by injectivity of α_i we have $(A, B) = (A', B')$. The same holds for β and γ .

Moreover, the orbits of $\alpha(A, B'), \beta(B, C'), \gamma(C, A')$ form a triangle in $\text{Sym}^k \mathcal{C}$ iff $A = A', B = B',$ and $C = C'$. For suppose there exist points X, Y, Z and permutations $\pi_1, \pi_2, \pi_3 \in S_k$ for which $(X, Y) \in R_I, (Y, Z) \in R_J$ and $(Z, X) \in R_K$, where $I = \pi_1 \alpha(A, B'), J = \pi_2 \beta(B, C'),$ and $K = \pi_3 \gamma(C, A')$. Then we must have $\pi_1 = \pi_2 = \pi_3$ because $\alpha_i(A_i, B'_i), \beta_j(B_j, C'_j), \gamma_k(C_k, A'_k)$ cannot form a triangle unless $i = j = k$, and then $A' = A, B = B'$ and $C = C'$ follow from the fact that these equalities hold for each coordinate (by properties of $\alpha_i, \beta_i, \gamma_i$). Thus $\text{Sym}^k \mathcal{C}$ realizes $\langle L, M, N \rangle$, as claimed. \square

COROLLARY 5.4. *Let \mathcal{C} be a commutative coherent configuration of rank r that realizes $\bigoplus_{i=1}^k \langle \ell_i, m_i, n_i \rangle$. Then symmetric powers of direct powers of \mathcal{C} prove the bound*

$$k \cdot \left(\left(\prod_{i=1}^k \ell_i m_i n_i \right)^{1/k} \right)^{\omega_s/3} \leq r.$$

More precisely, they come arbitrarily close to this bound.

Proof. By taking direct powers, \mathcal{C}^t realizes

$$\bigoplus_{I \in [k]^t} \left\langle \prod_{j \in [t]} \ell_{I_j}, \prod_{j \in [t]} m_{I_j}, \prod_{j \in [t]} n_{I_j} \right\rangle.$$

Setting $L = \prod_{i \in [k]} \ell_i$, $M = \prod_{i \in [k]} m_i$ and $N = \prod_{i \in [k]} n_i$, we have by Theorem 5.3 that $\text{Sym}^k \mathcal{C}^t$ realizes

$$\langle L^{tk^{t-1}}, M^{tk^{t-1}}, N^{tk^{t-1}} \rangle$$

and has rank $\binom{r+k^{t-1}}{k^{t-1}}$. By Proposition 3.5 we have

$$\begin{aligned} (LMN)^{\omega_s tk^{t-1}/3} &\leq \binom{r+k^{t-1}}{k^{t-1}} \\ &\leq \left(\frac{e(r+k^{t-1})}{k^{t-1}} \right)^{k^{t-1}} \\ &\leq \left(\frac{2er^t}{k^t} \right)^{k^t}, \end{aligned}$$

where the last inequality uses the fact that $k \leq r$. Taking tk^t -th roots and letting t go to infinity, we obtain the bound $(LMN)^{\omega_s/(3k)} \leq r/k$. \square

By weighting the independent matrix multiplications appropriately, we find that the geometric mean can be replaced by the arithmetic mean, to obtain a bound on ω_s identical to the asymptotic sum inequality (2.1):

THEOREM 5.5. *Let \mathcal{C} be a commutative coherent configuration of rank r that realizes $\bigoplus_{i=1}^k \langle \ell_i, m_i, n_i \rangle$. Then symmetric powers of direct powers of \mathcal{C} prove the bound $\sum_i (\ell_i m_i n_i)^{\omega_s/3} \leq r$.*

Proof. Fix an integer N and $\mu = (\mu_1, \dots, \mu_k)$ satisfying $\mu_i \geq 0$ and $\sum_i \mu_i = N$. Then the direct product \mathcal{C}^N realizes $L = \binom{N}{\mu}$ independent copies of $\langle \prod_i \ell_i^{\mu_i}, \prod_i m_i^{\mu_i}, \prod_i n_i^{\mu_i} \rangle$ (the key is that now these are all the same size). Applying Corollary 5.4, we find that symmetric powers of direct powers of \mathcal{C} prove the bound

$$(5.2) \quad \binom{N}{\mu} \prod_i (\ell_i m_i n_i)^{\mu_i \omega_s/3} \leq r^N.$$

Summing this inequality over all μ gives

$$\left(\sum_i (\ell_i m_i n_i)^{\omega_s/3} \right)^N \leq \binom{N+k-1}{k-1} \cdot r^N,$$

and the theorem follows by taking N -th roots and letting N go to infinity. Note that for each N , by an averaging argument, there must be a particular distribution μ for which the left hand side of (5.2) is at least $(\sum_i (\ell_i m_i n_i)^{\omega_s/3})^N / \binom{N+k-1}{k-1}$ and this is a concrete sequence of coherent configurations that prove the same bound in the limit. \square

The results of this section are not specific to coherent configurations. Given any algebra A , the analogous construction is to look at the subalgebra of $A^{\otimes n}$ invariant under the action of S_n .

Before using Theorem 5.5 to obtain bounds on ω , we briefly contrast other proofs of the asymptotic sum inequality with the above proof, which seems structurally different as we now explain. The standard proof of the asymptotic sum inequality takes a tensor T realizing $\bigoplus_{i=1}^k \langle \ell_i, m_i, n_i \rangle$ and finds $k!$ independent copies of $\langle \prod_i \ell_i, \prod_i m_i, \prod_i n_i \rangle$ in $T^{\otimes k}$. By performing block matrix multiplication, these are capable of realizing the *larger* matrix multiplication instance $\langle K \cdot \prod_i \ell_i, K \cdot \prod_i m_i, K \cdot \prod_i n_i \rangle$ (where $K \approx k!^{1/\omega}$), and the general bound follows after some manipulations analogous to our Corollary 5.4 and Theorem 5.5. In contrast, our proof finds a *single* copy of $\langle \prod_i \ell_i, \prod_i m_i, \prod_i n_i \rangle$ in $T^{\otimes k}$, and then uses the fact that T has special structure—it is the structural tensor of an algebra—to argue that the same matrix multiplication instance survives after symmetrizing the k -th power. Symmetrizing reduces the rank, and thus the s -rank actually *shrinks* enough to obtain the same bound. We know of no other proof that works by shrinking the rank (including the proof in [7] for independent matrix multiplications realized in group algebras).

5.1 Nontrivial bounds on ω Using Theorem 5.5, we can convert all of the results of [7] into realizations of a single matrix multiplication tensor in a *commutative* coherent configuration, namely a symmetric power of an abelian group. While the starting constructions are not new, the final algorithms are (they use machinery introduced in this paper). They establish that commutative coherent configurations suffice to prove nontrivial bounds on ω , and even point to a specific family of commutative coherent configurations that we conjecture is capable of proving $\omega = 2$.

THEOREM 5.6. *There exist commutative coherent configurations that prove s -rank exponent bounds $\omega_s \leq 2.48$, $\omega_s \leq 2.41$, and $\omega_s \leq 2.376$, and thus corresponding exponent bounds $\omega \leq 2.72$, $\omega \leq 2.62$, and $\omega \leq 2.564$, respectively.*

Proof. Apply Theorem 5.5 to the abelian group constructions of Proposition 3.8 in [7], Theorems 3.3 and 6.6 in [7], and the generalization matching [9] (as stated in [7] but not described in detail), respectively. Each of these constructions from [7], when viewing groups as coherent configurations and adopting the language of this paper, gives coherent configurations satisfying Definition 5.1. Apply Theorem 3.6 to the resulting s -rank exponent bounds to obtain the claimed bounds on ω . \square

The specific exponent bounds cited above of course all suffer from the 50% penalty introduced by Theorem 3.6. But the numbers themselves should not obscure the main point,

which is that matrix multiplication via coherent configurations is a viable approach to proving $\omega = 2$. Indeed, we conjecture that commutative coherent configurations are sufficient to prove $\omega = 2$:

CONJECTURE 5.7. *There exist commutative coherent configurations \mathcal{C}_n realizing $\langle n, n, n \rangle$ and of rank $n^{2+o(1)}$.*

Such a family of commutative coherent configurations would prove $\omega = 2$. If Conjecture 3.4 or 4.7 from [7] hold, then Conjecture 5.7 holds via Theorem 5.5. We note that recent work of Alon, Shpilka, and Umans [1] shows that Conjecture 3.4 from [7] contradicts a sunflower conjecture, although there is no strong consensus that this particular sunflower conjecture is true. Among the various combinatorial/algebraic conjectures implying $\omega = 2$, Conjecture 5.7 is the weakest (it is implied by the others), which makes it the “easiest” among these potential routes to proving $\omega = 2$.

6 Families of coherent configurations

In this section we discuss the suitability of broad classes of coherent configurations for proving bounds on ω .

6.1 Coherent configurations with many fibers By property (1) in the definition of a coherent configuration, there is a subset of the classes that form a partition of the diagonal, and we call these classes the *fibers* of the coherent configuration. We noted in Section 4.2 that coherent configurations with more than one fiber are noncommutative. More interestingly for our application, we will see shortly that n fibers suffice to embed $n \times n$ matrix multiplication. This observation generalizes Example 4.6.

The fibers of a coherent configuration \mathcal{C} correspond to a partition of the points into subsets $\mathcal{C}_1, \dots, \mathcal{C}_n$. Then it follows from property (3) in the definition that the classes of \mathcal{C} form a refinement of the subsets $\mathcal{C}_i \times \mathcal{C}_j$.

PROPOSITION 6.1. *Every coherent configuration with n fibers realizes $\langle n, n, n \rangle$.*

Proof. Let \mathcal{C} be a coherent configuration with n fibers and corresponding partition $\mathcal{C}_1, \dots, \mathcal{C}_n$, and let x_1, \dots, x_n be a system of distinct representatives for $\mathcal{C}_1, \dots, \mathcal{C}_n$. Define $\alpha(a, b)$ to be the class of \mathcal{C} containing (x_a, x_b) , $\beta(b, c)$ to be the class containing (x_b, x_c) and $\gamma(c, a)$ to be the class containing (x_c, x_a) . It is easy to verify that these functions satisfy Definition 4.1. \square

However, this generic embedding does not lead to nontrivial bounds on ω_s , because a similar argument shows that one of the character degrees must be at least n . Let \mathcal{C}' be the coherent configuration with the same points as \mathcal{C} and $\mathcal{C}_i \times \mathcal{C}_j$ as its classes. Then the adjacency algebra of \mathcal{C}' is a subalgebra of that of \mathcal{C} (the adjacency matrix for $\mathcal{C}_i \times \mathcal{C}_j$ is the sum of the adjacency matrices for the classes of \mathcal{C}

contained in $\mathcal{C}_i \times \mathcal{C}_j$), and it is not hard to check that the adjacency algebra of \mathcal{C}^ℓ is isomorphic to $\mathbb{C}^{n \times n}$. However, $\mathbb{C}^{n \times n}$ cannot be isomorphic to a subalgebra of a semisimple algebra $\mathbb{C}^{d_1 \times d_1} \times \dots \times \mathbb{C}^{d_k \times d_k}$ unless $d_i \geq n$ for some i . To see why, note that projection onto the factors in the semisimple algebra would yield representations of dimension d_i for $\mathbb{C}^{n \times n}$. Because $\mathbb{C}^{n \times n}$ is a simple algebra, these projections must vanish unless $d_i \geq n$.

6.2 Schurian coherent configurations Recall that the *Schurian* coherent configurations are those obtained from actions of groups on sets, and that such a coherent configuration is an association scheme iff the action is transitive.

In the special case of a finite group acting on itself by right multiplication, our framework is equivalent to the triple product property from [8]. Thus, the conjectures of [7], which all imply $\omega = 2$ via the triple product property in groups, imply that Schurian coherent configurations are sufficient to achieve $\omega_s = 2$.

More interestingly, observe that the coherent configurations arising in Theorem 5.6 and Conjecture 5.7 are in fact commutative Schurian coherent configurations. This is because symmetric powers of coherent configurations arising from abelian groups, which are commutative, are Schurian via a wreath product action: if G is a group and \mathcal{C} is the associated coherent configuration, then $\text{Sym}^k \mathcal{C}$ is the Schurian coherent configuration arising from $S_k \times G^k$ acting on G^k .

Thus commutative Schurian coherent configurations, which arise from transitive group actions, already prove nontrivial bounds on ω_s (and ω), and if either of the two conjectures in [7] is true, then they suffice to prove $\omega_s = 2$.

6.3 Group association schemes Another generic way to obtain a Schurian coherent configuration is to consider $G \times G$ acting on G via $(x, y) \cdot g = xgy^{-1}$. This gives rise to a commutative coherent configuration (regardless of whether G is commutative or not, which is attractive for our application) called the *group association scheme*, whose classes are identified with conjugacy classes of G (i.e., class $R_i = \{(g, h) : gh^{-1} \in C_i\}$, where C_i is the i -th conjugacy class). Here we show that group association schemes suffice to prove nontrivial bounds on ω_s , and that if either of the two conjectures in [7] is true, then group association schemes suffice to prove $\omega_s = 2$.

We need the following definition from [7] (Definition 5.1):

DEFINITION 6.2. We say that n triples of subsets A_i, B_i, C_i of a group H satisfy the *simultaneous triple product property* if

$$\begin{aligned} a_i^{-1} a'_j b_j^{-1} b'_k c_k^{-1} c'_i &= 1 \\ \Leftrightarrow \\ i = j = k \text{ and } a_i &= a'_j, b_j = b'_k, c_k = c'_i \end{aligned}$$

holds for all i, j, k and $a_i \in A_i, a'_j \in A_j, b_j \in B_j, b'_k \in B_k, c_k \in C_k, c'_i \in C_i$.

When this holds, the coherent configuration associated with the right action of H on itself realizes $\oplus_i \langle |A_i|, |B_i|, |C_i| \rangle$ via functions $\alpha_i, \beta_i,$ and γ_i defined on $A_i \times B_i, B_i \times C_i,$ and $C_i \times A_i,$ respectively. Specifically, $\alpha_i(a, b)$ is the class containing the pair (a, b) , etc. Then Definition 5.1 amounts to the simultaneous triple product property.

The paper [7] describes the following constructions, among others:

- (1) It follows from Theorem 3.3 and Section 6.3 in [7] that for all $m > 2$, and ℓ sufficiently large, there are n triples A_i, B_i, C_i of subsets of $(\mathbb{Z}/m\mathbb{Z})^{3\ell}$ satisfying the simultaneous triple product property with $n = (27/4)^{\ell - o(\ell)}$ and $|A_i||B_i||C_i| = (m-2)^{3\ell}$ for all i . Applying Theorem 5.5 and taking the limit as $\ell \rightarrow \infty$ yields

$$\omega_s \leq \frac{3 \log m - \log(27/4)}{\log(m-2)},$$

which is optimized by $m = 10$ (giving $\omega_s \leq 2.41$).

- (2) Either of the two conjectures in [7] implies the existence of subsets satisfying the simultaneous triple product property in an abelian group H with

$$|A_i| = |B_i| = |C_i| = t \geq n^\epsilon$$

for $1 \leq i \leq n$ and $|H| = (t^2 n)^{1+o(1)}$ (as $n \rightarrow \infty$ with $\epsilon > 0$ fixed), which would prove $\omega = 2$.

THEOREM 6.3. Let H be an abelian group, and suppose n triples of subsets A_i, B_i, C_i in H satisfy the simultaneous triple product property. Let $G = S_n \times H^n$, and define

$$\begin{aligned} A &= A_1 \times A_2 \times \dots \times A_n \\ B &= B_1 \times B_2 \times \dots \times B_n \\ C &= C_1 \times C_2 \times \dots \times C_n, \end{aligned}$$

viewed as subsets of G via the natural embedding of H^n in G . Let \mathcal{C} be the group association scheme of G . Then the subsets A, B, C satisfy the requirements of Proposition 4.7 with respect to \mathcal{C} (i.e., for the action of $G \times G$ on G), so \mathcal{C} realizes $\langle |A|, |B|, |C| \rangle$.

Proof. We will write elements of G as $h\pi$, with $h \in H^n, \pi \in S_n$, and we will use $\pi \cdot h$ to denote the permutation action of S_n on H^n . The semidirect product satisfies $\pi h = (\pi \cdot h)\pi$ for $\pi \in S_n$ and $h \in H^n$.

Suppose we have $f = (f_1, f_2), g = (g_1, g_2), h = (h_1, h_2)$ in $G \times G$ and $a, a' \in A, b, b' \in B, c, c' \in C$ for which

$$(6.1) \quad \begin{aligned} fgh &= 1 \\ f_1 a f_2^{-1} &= a' \\ g_1 b g_2^{-1} &= b' \\ h_1 c h_2^{-1} &= c'. \end{aligned}$$

We wish to conclude that $a = a', b = b', c = c'$. From the latter three equations in (6.1), we see that $f_1 = x_1\pi, f_2 = x_2\pi$ for some $x_1, x_2 \in H^n$ and $\pi \in S_n$. Similarly, $g_1 = y_1\rho, g_2 = y_2\rho$ and $h_1 = z_1\tau, h_2 = z_2\tau$. Now, using the commutativity of H , the three equations become

$$(6.2) \quad \begin{aligned} x_1x_2^{-1} &= a'(\pi \cdot a^{-1}) \\ y_1y_2^{-1} &= b'(\rho \cdot b^{-1}) \\ z_1z_2^{-1} &= c'(\tau \cdot c^{-1}). \end{aligned}$$

From $fgh = 1$ we have $f_1g_1h_1 = 1$, which implies

$$1 = f_1g_1h_1 = x_1(\pi \cdot y_1)((\pi\rho) \cdot z_1).$$

Similarly, from $fgh = 1$ we have $f_2g_2h_2 = 1$ and hence

$$1 = f_2g_2h_2 = x_2(\pi \cdot y_2)((\pi\rho) \cdot z_2).$$

Using the commutativity of H , we obtain from these two equations

$$x_1x_2^{-1}(\pi \cdot (y_1y_2^{-1}))((\pi\rho) \cdot (z_1z_2^{-1})) = 1,$$

which combined with (6.2) yields

$$a'(\pi \cdot a^{-1})(\pi \cdot b')((\pi\rho) \cdot b^{-1})((\pi\rho) \cdot c')((\pi\rho\tau) \cdot c^{-1}) = 1.$$

Now, $fgh = 1$ implies $\pi\rho\tau = 1$, so we obtain

$$(\pi \cdot (a^{-1}b'))((\pi\rho) \cdot (b^{-1}c'))(c^{-1}a') = 1.$$

This implies via the simultaneous triple product property that $\pi = \pi\rho = 1$. We conclude that $\pi = \rho = \tau = 1$, and then that $a = a', b = b', c = c'$ as desired. \square

To determine what bounds on ω_s can be expected, we need to know the rank of this group association scheme, i.e., the number of conjugacy classes in $S_n \times H^n$:

LEMMA 6.4. *There is a constant C such that for every abelian group H , if $n \leq |H|$ then the number of conjugacy classes of $S_n \times H^n$ is at most $C^n |H|^n / n^n$.*

This is a crude bound, but it will suffice for our purposes.

Proof. It is not difficult to prove the following description of the conjugacy classes in the group $S_n \times H^n$. Each element of this group can be written as $h\pi$ with $h \in H^n$ and $\pi \in S_n$. The cycle type of π is preserved under conjugation, and the sum of the elements of H in the coordinates of h corresponding to each cycle of π is also preserved. Furthermore, these invariants completely specify the conjugacy class. Thus, each conjugacy class is specified by a multiset of pairs consisting of a cycle length and an element of H , where the cycle lengths must sum to n .

The possible cycle types correspond to partitions of n , and the number of them grows subexponentially as $n \rightarrow \infty$.

More elementarily, there are 2^{n-1} compositions of n (ways of writing n as an ordered sum of positive integers), and therefore at most 2^{n-1} partitions of n .

Suppose the permutation has c_i cycles of length i , with $\sum_{i=1}^n ic_i = n$. Then there are

$$\prod_{i=1}^n \binom{|H| + c_i - 1}{c_i}$$

ways to choose the elements of H corresponding to these cycles. Thus, bounding the number of conjugacy classes in G amounts to bounding how large this product can be.

We have

$$\begin{aligned} \prod_{i=1}^n \binom{|H| + c_i - 1}{c_i} &\leq \prod_{i=1}^n \left(\frac{e(|H| + c_i - 1)}{c_i} \right)^{c_i} \\ &\leq (2e)^n \prod_{i=1}^n \frac{|H|^{c_i}}{c_i^{c_i}} \\ &\leq (2e)^n \prod_{i=1}^n \frac{|H|^{ic_i}}{c_i^{c_i} n^{(i-1)c_i}} \\ &\leq (2e)^n \frac{|H|^n}{n^n} \prod_{i=1}^n \left(\frac{n}{c_i} \right)^{c_i}. \end{aligned}$$

If we set $x_i = c_i/n$, then

$$\prod_{i=1}^n \left(\frac{n}{c_i} \right)^{c_i} = e^{-n \sum_{i=1}^n x_i \log x_i},$$

where \log denotes the natural logarithm. Thus, to complete the proof we must show that $-\sum_{i=1}^n x_i \log x_i$ is bounded independently of n , whenever $x_i \geq 0$ and $\sum_{i=1}^n ix_i = 1$. The maximum can be found using Lagrange multipliers. One must deal with the boundary cases when $x_i = 0$ for some i , and we provide the details below.

Suppose x_1, \dots, x_n maximize $-\sum_{i=1}^n x_i \log x_i$ subject to $\sum_{i=1}^n ix_i = 1$ and $x_i \geq 0$. The desired result is trivial when only one of x_1, \dots, x_n is nonzero. Otherwise, let z_1, \dots, z_m be the nonzero elements among x_1, \dots, x_n . The equation $\sum_{i=1}^n ix_i = 1$ becomes $\sum_{i=1}^m y_i z_i = 1$, where $y_1 < y_2 < \dots < y_m$ are positive integers. Then there is a Lagrange multiplier λ such that $-1 - \log z_i = \lambda y_i$ for all i , and hence

$$-\sum_{i=1}^m z_i \log z_i = \sum_{i=1}^m z_i (\lambda y_i + 1) = \lambda + \sum_{i=1}^m z_i \leq \lambda + 1.$$

To bound λ , note that $z_i = e^{-1-\lambda y_i}$ and hence

$$\sum_{i=1}^m y_i e^{-1-\lambda y_i} = 1,$$

while for $\lambda > 1$ we have

$$\sum_{i=1}^m y_i e^{-1-\lambda y_i} < \sum_{j=1}^{\infty} j e^{-1-j} < 1.$$

Thus, $\lambda \leq 1$ and so $-\sum_{i=1}^m z_i \log z_i \leq 2$. Combining the estimates so far shows that we can take $C = 4e^3$ in the lemma statement. (The best possible constant is of course much smaller.) \square

This bound on the rank of these group association schemes is precisely what is needed to recover the desired bounds from the simultaneous triple product property constructions listed earlier in this section. E.g., applying Theorem 6.3 and Lemma 6.4 to the second example yields

$$t^{n\omega_s} \leq C^n \frac{(t^2 n)^{n(1+o(1))}}{n^n},$$

which is equivalent to

$$\omega_s \log t \leq \log C + (2 + o(1)) \log t + o(\log n).$$

Because $t \geq n^\epsilon$, we get $\omega_s = 2$ in the limit as $n \rightarrow \infty$.

COROLLARY 6.5. *There exist group association schemes that prove $\omega_s \leq 2.41$ (and hence $\omega \leq 2.62$). If either of the conjectures in [7] is true, then there are group association schemes that prove $\omega_s = 2$ (and hence $\omega = 2$).*

More generally, one can imitate the transition from Corollary 5.4 (which is analogous to Theorem 6.3) to Theorem 5.5 to give a proof of Theorem 5.5 from [7] using group association schemes.

References

- [1] N. Alon, A. Shpilka, and C. Umans. On sunflowers and matrix multiplication. Proceedings of the 27th IEEE Conference on Computational Complexity, 26–29 June 2012, Porto, Portugal, IEEE Computer Society, pp. 214–223.
- [2] D. Bini. Relations between exact and approximate bilinear algorithms. Applications. *Calcolo*, 17:87–97, 1980.
- [3] D. Bini, M. Capovani, G. Lotti, and F. Romani. $O(n^{2.7799})$ complexity for matrix multiplication. *Inf. Proc. Letters*, 8:234–235, 1979.
- [4] M. Bläser. Complexity of bilinear algorithms (course notes taken by F. Bendun). <http://www-cc.cs.uni-saarland.de/media/oldmaterial/bc.pdf>, 2009.
- [5] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic Complexity Theory*, volume 315 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, 1997.
- [6] P. J. Cameron. *Permutation Groups*, volume 45 of *London Mathematical Society Students Texts*. Cambridge University Press, 1999.
- [7] H. Cohn, R. Kleinberg, B. Szegedy, and C. Umans. Group-theoretic algorithms for matrix multiplication. Proceedings of the 46th Annual Symposium on Foundations of Computer Science, 23–25 October 2005, Pittsburgh, PA, IEEE Computer Society, pp. 379–388, [arXiv:math.GR/0511460](https://arxiv.org/abs/math/0511460).
- [8] H. Cohn and C. Umans. A group-theoretic approach to fast matrix multiplication. Proceedings of the 44th Annual Symposium on Foundations of Computer Science, 11–14 October 2003, Cambridge, MA, IEEE Computer Society, pp. 438–449, [arXiv:math.GR/0307321](https://arxiv.org/abs/math/0307321).
- [9] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9:251–280, 1990.
- [10] A. M. Davie and A. J. Stothers. Improved bound for complexity of matrix multiplication. *Proceedings of the Royal Society of Edinburgh, Section A*. To appear.
- [11] J. Håstad. Tensor rank is NP-Complete. *J. Algorithms*, 11:644–654, 1990.
- [12] D. G. Higman. Coherent configurations I. *Rend. Sem. Mat. Univ. Padova*, 44:1–25, 1970.
- [13] D. G. Higman. Combinatorial considerations about permutation groups. Lectures given in 1971. Mathematical Institute, Oxford University, Oxford, 1972.
- [14] D. G. Higman. Coherent configurations I. Ordinary representation theory. *Geometriae Dedicata*, 4:1–32, 1975.
- [15] R. Salem and D. C. Spencer. On sets of integers which contain no three terms in arithmetical progression. *Proc. Nat. Acad. Sci. USA*, 28:561–563, 1942.
- [16] A. Schönhage. Partial and total matrix multiplication. *SIAM J. Comp.*, 10:434–455, 1981.
- [17] A. Stothers. On the complexity of matrix multiplication. Ph.D. dissertation, University of Edinburgh, 2010.
- [18] V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13:354–356, 1969.
- [19] V. Strassen. Relative bilinear complexity and matrix multiplication. *J. Reine Angew. Math.*, 375/376:406–443, 1987.
- [20] V. Vassilevska Williams. Multiplying matrices faster than Coppersmith-Winograd. Proceedings of the 44th ACM Symposium on Theory of Computing, 19–22 May 2012, New York, NY, Association for Computing Machinery, pp. 887–898.