Group-theoretic Algorithms for Matrix Multiplication

Henry Cohn* Robert Kleinberg† Balázs Szegedy‡ Christopher Umans§

Abstract

We further develop the group-theoretic approach to fast matrix multiplication introduced by Cohn and Umans, and for the first time use it to derive algorithms asymptotically faster than the standard algorithm. We describe several families of wreath product groups that achieve matrix multiplication exponent less than 3*, the asymptotically fastest of which achieves exponent* 2.41*. We present two conjectures regarding specific improvements, one combinatorial and the other algebraic. Either one would imply that the exponent of matrix multiplication is* 2*.*

1. Introduction

The task of multiplying matrices is one of the most fundamental problems in algorithmic linear algebra. Matrix multiplication itself is a important operation, and its importance is magnified by the number of similar problems that are reducible to it.

Following Strassen's discovery [9] of an algorithm for $n \times n$ matrix multiplication in $O(n^{2.81})$ operations, a sequence of improvements has achieved ever better bounds on the *exponent of matrix multiplication*, which is the smallest real number ω for which $n \times n$ matrix multiplication can be performed in $O(n^{\omega+\varepsilon})$ operations for each $\varepsilon > 0$. The asymptotically fastest algorithm known is due to Coppersmith and Winograd [3], and it proves that $\omega < 2.376$. Since 1990, there have been no better upper bounds proved on ω, although it is widely believed that $\omega = 2$.

Recently, Cohn and Umans [2] proposed a new group-theoretic approach to devising matrix multiplication algorithms. In this framework, one selects a finite group G satisfying a certain property that allows $n \times n$ matrix multiplication to be reduced to multiplication of elements of the *group algebra* $\mathbb{C}[G]$. This latter multiplication is performed via a Fourier transform, which reduces it to several smaller matrix multiplications, whose sizes are the *character degrees* of G. This naturally gives rise to a recursive algorithm whose running time depends on the character degrees. Thus the problem of devising matrix multiplication algorithms in this framework is imported into the domain of group theory and representation theory.

One of the main contributions of [2] was to demonstrate that several diverse families of non-abelian groups support the reduction of $n \times n$ matrix multiplication to group algebra multiplication. These include, in particular, families of groups of size $n^{2+o(1)}$. The existence of such families is a necessary condition for the group-theoretic approach to prove $\omega = 2$, although it is not sufficient.

The main question raised in [2] is whether the proposed approach could prove nontrivial bounds on ω, i.e., prove $\omega < 3$. This was shown to be equivalent to a question in representation theory, Question 4.1 in [2]: is there a group G with subsets S_1, S_2, S_3 that satisfy the *triple product property* (see Definition 1.3 below), and for which $|S_1||S_2||S_3| > \sum d_i^3$, where $\{d_i\}$ is the set of character degrees of G?

In this paper we resolve this question in the affirmative, which immediately gives a simple matrix multiplication algorithm in the group-theoretic framework that has running time $O(n^{2.9088})$. The group we construct for this purpose is a *wreath product*, and in subsequent sections we describe similar constructions that produce algorithms with running times $O(n^{2.48})$ and $O(n^{2.41})$.

The main challenge in each case is to describe the three subsets of the group that satisfy the triple product property. We give two ways of organizing these descriptions, both of which give rise to the $O(n^{2.48})$ algorithm relatively simply. We also advance two natural conjectures related to these formulations, each of which would imply that $\omega = 2$. The first is a combinatorial conjecture (Conjecture 3.4), and the second is an algebraic conjecture (Conjecture 4.7).

The three subsets underlying the $O(n^{2.41})$ algorithm are described in terms of a combinatorial object we call a *Uniquely Solvable Puzzle* (or USP), which is a weakening of the combinatorial object in our first conjecture. An *op-*

[*] Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399, cohn@microsoft.com

[†] Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, and Department of Computer Science, Cornell University, Ithaca, NY 14853, rdk@cs.cornell.edu

[‡] Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399, szegedy@microsoft.com

[§] Department of Computer Science, California Institute of Technology, Pasadena, CA 91125, umans@cs.caltech.edu; supported in part by NSF grant CCF-0346991 and an Alfred P. Sloan Research Fellowship.

timal USP construction can be extracted from Coppersmith and Winograd's paper [3].

In fact, the reader familiar with Strassen's 1987 paper [10] and Coppersmith and Winograd's paper [3] (or the presentation of this material in, for example, [1]) will recognize that our exponent bounds of 2.48 and 2.41 match bounds derived in those works. It turns out that with some effort the algorithms in [10] and [3], including Coppersmith and Winograd's $O(n^{2.376})$ algorithm, all have analogues in our group-theoretic framework. The translation does not appear to be systematic: the algorithms are based on similar principles, but in fact they are not identical (the actual operations performed on matrix entries do not directly correspond), and we know of no group-theoretic interpretations of any earlier algorithms. We defer a complete account of this connection to the full version of this paper.

We believe that, compared to existing algorithms, our group-theoretic algorithms are simpler to state and simpler to analyze. They are situated in a clearer conceptual and mathematical framework, in which, for example, the two conjectures mentioned above are natural and easy to identify. Finally, they avoid various complications of earlier algorithms. For example, they substitute the discrete Fourier transform, together with some elementary facts in representation theory, for the seemingly ad hoc trilinear form identities in introduced in [10], and they completely avoid the need to deal with degenerations and border rank of tensors.

1.1. Outline

In the rest of this section, we establish notation and review background from [2] on the group-theoretic approach to fast matrix multiplication. Section 2 describes the simplest group we have found that can prove a nontrivial bound on the exponent of matrix multiplication. In Sections 3 and 4, we carry out a more elaborate construction in two different ways, each of which has the potential of reaching $\omega = 2$ although $\omega < 2.48$ is the best we can achieve so far by these methods. The most fundamental conceptual contribution in this paper is the *simultaneous triple product property*, which we introduce in Section 5. It extends the triple product property from [2], and it encompasses and illuminates all of our other constructions, as we explain in Section 6. Finally, in Section 7 we show that any bound provable via the simultaneous triple product property can in fact be proved using only the approach of [2].

1.2. Preliminaries and notation

As usual ω denotes the exponent of matrix multiplication over $\mathbb{C}$.

The set $\{1, 2, \dots, k\}$ is denoted $[k]$. We write $A \setminus B = \{a \in A : a \notin B\}$ and if A and B are subsets of an abelian group we set $A - B = \{a - b : a \in A, b \in B\}$.

The cyclic group of order k is denoted Cyc_k (with additive notation for the group law), and the symmetric group on a set S is denoted $\mathrm{Sym}(S)$ (or Sym_n instead of $\mathrm{Sym}([n])$). If G is a group and R is a ring, then $R[G]$ will denote the group algebra of G with coefficients in R.

When we discuss a group action, it will always be a left action unless otherwise specified. If G and H are groups with a left action of G on H (where the action of g on h is written $g \cdot h$), then the semidirect product $H \rtimes G$ is the set $H \times G$ with the multiplication law

$$(h_1, g_1), (h_2, g_2) = (h_1(g_1 \cdot h_2), g_1 g_2).$$

We almost always identify H with the subset $H \times \{1\}$ and G with $\{1\} \times G$, so that (h, g) simply becomes the product hg.

For a right action of G on H, with the action of g on h written h^g, the semidirect product $G \ltimes H$ is $G \times H$ with the multiplication law

$$(g_1, h_1)(g_2, h_2) = (g_1 g_2, h_1^{g_2} h_2).$$

As in the previous case we identify G and H with the corresponding subsets of $G \ltimes H$.

Other than for Lemma 1.2, which is not required for the main results of this paper, we will use only the following basic facts from representation theory. The group algebra $\mathbb{C}[G]$ of a finite group G decomposes as the direct product

$$\mathbb{C}[G] \cong \mathbb{C}^{d_1 \times d_1} \times \cdots \times \mathbb{C}^{d_k \times d_k}$$

of matrix algebras of orders $d_1, \dots, d_k$. These orders are the character degrees of G, or the dimensions of the irreducible representations. It follows from computing the dimensions of both sides that $|G| = \sum_i d_i^2$. It is also easy to prove that if G has an abelian subgroup A, then all the character degrees of G are less than or equal to the index $[G : A]$ (Proposition 2.6 in [5]) . See [6] and [5] for further background on representation theory.

The following elementary lemma will prove useful several times:

Lemma 1.1. *Let $s_1, s_2, \dots, s_n$ be nonnegative real numbers, and suppose that for every vector $\mu = (\mu_1, \dots, \mu_n)$ of nonnegative integers for which $\sum_{i=1}^n \mu_i = N$ we have*

$$\binom{N}{\mu} \prod_{i=1}^{n} s_i^{\mu_i} \le C^N.$$

Then $\sum_{i=1}^n s_i \le C$.

Proof. For each probability distribution $p = (p_1, \dots, p_n)$, we can let N tend to infinity and choose μ so that

$\lim_{N\to\infty} \mu/N = p$. As $N \to \infty$, the inequality in the hypothesis of the lemma yields

$$-\sum_i p_i \log p_i + \sum_i p_i \log s_i \le \log C$$

after taking the N-th root and the logarithm. Setting $S = \sum_i s_i$ and $p_i = s_i/S$ proves $\log S \le \log C$, as desired. □

Occasionally we will need to bound the character degrees of wreath products, but the proof can be skipped by readers not comfortable with representation theory:

Lemma 1.2. *Let $\{d_k\}$ be the character degrees of finite group H and let $\{c_j\}$ be the character degrees of* $\mathrm{Sym}_n \ltimes H^n$ *(where* Sym_n *acts by permuting the coordinates). Then*

$$\sum_j c_j^\omega \le (n!)^{\omega-1}\left(\sum_k d_k^\omega\right)^n.$$

Proof. When H is abelian, we can use the elementary facts that the character degrees of $\mathrm{Sym}_n \ltimes H^n$ are at most $n!$ (which is the index of H^n in $\mathrm{Sym}_n \ltimes H^n$) and that $\sum_j c_j^2 = |\mathrm{Sym}_n \ltimes H^n|$ to obtain

$$\sum_j c_j^\omega \le (n!)^{\omega-2}\sum_j c_j^2 = (n!)^{\omega-1}|H|^n.$$

For general H, we need further information regarding the character degrees of $\mathrm{Sym}_n \ltimes H^n$. From Theorem 25.6 in [5] we get the following description: The symmetric group Sym_n acts on the irreducible representations of H^n by permuting the n factors. Let V be any irreducible representation of H^n, and let $G_V \subseteq \mathrm{Sym}_n$ be the subgroup that fixes V (in the action on such representations). Then V extends to a representation of $G_V \ltimes H^n$. Taking the tensor product with an irreducible representation W of G_V (with H^n acting trivially on W) and inducing to $\mathrm{Sym}_n \ltimes H^n$ yields an irreducible representation

$$\mathrm{Ind}_{G_V \ltimes H^n}^{\mathrm{Sym}_n \ltimes H^n}(W \otimes_{\mathbb{C}} V)$$

of $\mathrm{Sym}_n \ltimes H^n$. All irreducible representations of $\mathrm{Sym}_n \ltimes H^n$ arise in this way. Two such representations are isomorphic iff the choices of W are isomorphic and the choices of V are equivalent under the action of Sym_n on irreducible representations of H^n. In other words, if we look at all choices of V and W, each representation is counted $n!/|G_V|$ times (because that is the size of V's orbit under Sym_n).

The dimension of this representation is $(n!/|G_V|)\dim(W)\dim(V)$. Thus, because $\dim(W) \le |G_V|$ and $\sum_W \dim(W)^2 = |G_V|$,

$$\begin{aligned}\sum_j c_j^\omega &= \sum_V \frac{|G_V|}{n!}\sum_W\left(\frac{n!}{|G_V|}\dim(W)\dim(V)\right)^\omega\\ &\le (n!)^{\omega-1}\sum_V \dim(V)^\omega = (n!)^{\omega-1}\left(\sum_k d_k^\omega\right)^n,\end{aligned}$$

as desired. □

1.3. Background

In this subsection we summarize the necessary definitions and results from [2].

If S is a subset of a group, let $Q(S)$ denote the right quotient set of S, i.e., $Q(S) = \{s_1 s_2^{-1} : s_1, s_2 \in S\}$.

Definition 1.3 ([2]). *A group G* realizes *$\langle n_1, n_2, n_3\rangle$ if there are subsets $S_1, S_2, S_3 \subseteq G$ such that $|S_i| = n_i$, and for $q_i \in Q(S_i)$, if*

$$q_1 q_2 q_3 = 1$$

then $q_1 = q_2 = q_3 = 1$. We call this condition on S_1, S_2, S_3 the triple product property.

Lemma 1.4 ([2]). *If G realizes $\langle n_1, n_2, n_3\rangle$, then it does so for every permutation of n_1, n_2, n_3.*

Lemma 1.5 ([2]). *If $S_1, S_2, S_3 \subseteq G$ and $S_1', S_2', S_3' \subseteq G'$ satisfy the triple product property, then so do the subsets $S_1 \times S_1', S_2 \times S_2', S_3 \times S_3' \subseteq G \times G'$.*

Theorem 1.6 ([2]). *Let R be any algebra over $\mathbb{C}$ (not necessarily commutative). If G realizes $\langle n, m, p\rangle$, then the number of ring operations required to multiply $n \times m$ with $m \times p$ matrices over R is at most the number of operations required to multiply two elements of $R[G]$.*

One particularly useful construction from [2] involves permutations of the points in a triangular array. Let

$$\Delta_n = \{(a, b, c) \in \mathbb{Z}^3 : a + b + c = n - 1 \text{ and } a, b, c \ge 0\}.$$

Geometrically, these triples are barycentric coordinates for a triangular array of points with n points along each side, but it is more convenient to manipulate them algebraically.

For $x \in \Delta_n$, we write $x = (x_1, x_2, x_3)$. Let H_1, H_2, and H_3 be the subgroups of $\mathrm{Sym}(\Delta_n)$ that preserve the first, second, and third coordinates, respectively. Specifically,

$$H_i = \{\pi \in \mathrm{Sym}(\Delta_n) : (\pi(x))_i = x_i \text{ for all } x \in \Delta_n\}.$$

Theorem 1.7 ([2]). *The subgroups H_1, H_2, H_3 defined above satisfy the triple product property.*

Theorem 1.8 ([2]). *Suppose G realizes $\langle n, m, p\rangle$ and the character degrees of G are $\{d_i\}$. Then*

$$(nmp)^{\omega/3} \le \sum_i d_i^\omega.$$

Combining Theorem 1.8 with the fact that $\sum_i d_i^2 = |G|$ yields the following corollary, which is generally how the theorem is applied:

Corollary 1.9 ([2]). *Suppose G realizes $\langle n, m, p\rangle$ and has largest character degree d. Then $(nmp)^{\omega/3} \le d^{\omega-2}|G|$.*

2. Beating the sum of the cubes

Suppose G realizes $\langle n, m, p\rangle$ and has character degrees $\{d_i\}$. Theorem 1.8 yields a nontrivial bound on ω (by ruling out the possibility of $\omega = 3$) if and only if

$$nmp > \sum_i d_i^3.$$

Question 4.1 in [2] asks whether such a group exists. In this section we construct one, which shows that our methods do indeed prove nontrivial bounds on ω. The rest of the paper is logically independent of this example, but it serves as motivation for later constructions.

We do not know of any construction that makes use of small groups. We have used the computer program GAP [4] to verify by brute force search that no group of order less than 128 proves a nontrivial bound on ω using three subgroups (as opposed to subsets). Thus a construction must involve either fairly sizable groups or subsets other than subgroups, and in fact all of our constructions involve both.

The example in this section realizes matrix multiplication through subsets other than subgroups. However, the subsets are close to subgroups in the sense that they can be obtained from subgroups by deleting a small number of elements.

Let $H = \mathrm{Cyc}_n^3$, and let $G = H^2 \rtimes \mathrm{Cyc}_2$, where Cyc_2 acts on H^2 by switching the two factors of H. Let z denote the generator of Cyc_2. We write elements of G in the form $(a, b)z^i$, with $a, b \in H$ and $i \in \{0, 1\}$. Note that $z(a, b)z = (b, a)$.

Let H_1, H_2, H_3 be the three factors of Cyc_n in the product $H = \mathrm{Cyc}_n^3$, viewed as subgroups of H. For notational convenience, let $H_4 = H_1$. Define subsets $S_1, S_2, S_3 \subseteq G$ by

$$S_i = \{(a, b)z^j : a \in H_i \setminus \{0\}, b \in H_{i+1}, j \in \{0, 1\}\}.$$

We will prove in Lemma 2.1 that these subsets satisfy the triple product property.

To analyze this construction we need very little representation-theoretic information. The character degrees of G are all at most 2, because H^2 is an abelian subgroup of index 2. Then since the sum of the squares of the character degrees is $|G|$, the sum of their cubes is at most $2|G|$, which equals $4n^6$.

On the other hand, $|S_i| = 2n(n-1)$, so $|S_1||S_2||S_3| = 8n^3(n-1)^3$. For $n \geq 5$, this product is larger than $4n^6$. By Corollary 1.9, $(2n(n-1))^\omega \leq 2^{\omega-2}2n^6$. The best bound on ω is achieved by setting $n = 17$, in which case we obtain $\omega < 2.9088$.

It is a straightforward calculation in representation theory to determine how many of the character degrees are 2 (and how many are 1). That can be used to improve this analysis, but the bound on ω changes by less than 10^{-4}, so we do not present the details here.

All that remains is to prove the triple product property:

Lemma 2.1. *S_1, S_2, and S_3 satisfy the triple product property.*

Proof. Consider the triple product $q_1q_2q_3$ with $q_i \in Q(S_i)$, and suppose it equals the identity. Each quotient q_i is either of the form $(a_i, b_i)(-a_i', -b_i')$ or of the form $(a_i, b_i)z(-a_i', -b_i')$, with $a_i, a_i' \in H_i$ and $b_i, b_i' \in H_{i+1}$. There must be an even number of factors of z among the three elements q_1, q_2, q_3.

First, suppose there are none. We can write $q_1q_2q_3$ as

$$(a_1, b_1)(-a_1', -b_1')(a_2, b_2)(-a_2', -b_2')(a_3, b_3)(-a_3', -b_3'),$$

where $a_i, a_i' \in H_i$ and $b_i, b_i' \in H_{i+1}$. The product is thus equal to

$$(a_1 - a_1' + a_2 - a_2' + a_3 - a_3', b_1 - b_1' + b_2 - b_2' + b_3 - b_3'),$$

which is the identity iff $q_1 = q_2 = q_3 = 1$, since the triple product property holds (trivially) for H_1, H_2, H_3 in H.

Second, suppose two of q_1, q_2, q_3 contain a z. The product $q_1q_2q_3$ can be simplified as above to yield a sum in each coordinate, except now a_i and a_i' contribute to different coordinates when q_i contains a z, as do b_i and b_i'. There are thus two i's such that a_i and a_i' contribute to different coordinates. For one of those two i's, b_{i-1} and b_{i-1}' contribute to the same coordinate (where we interpret the subscripts modulo 3). The sum in the other coordinate contains one of a_i and a_i' but neither of b_{i-1} and b_{i-1}', and thus only one summand from H_i (because for each j, $a_j, a_j' \in H_j$ and $b_j, b_j' \in H_{j+1}$). Since a_i and a_i' are nonzero by the definition of S_i, the product $q_1q_2q_3$ cannot be the identity. □

3. Uniquely solvable puzzles

In this section we define a combinatorial object called a *strong USP*, which gives rise to a systematic construction of sets satisfying the triple product property in a wreath product. Using strong USPs we achieve $\omega < 2.48$, and we conjecture that there exist strong USPs that prove $\omega = 2$.

3.1. USPs and strong USPs

A *uniquely solvable puzzle* (USP) of width k is a subset $U \subseteq \{1, 2, 3\}^k$ satisfying the following property:

> For all permutations $\pi_1, \pi_2, \pi_3 \in \mathrm{Sym}(U)$, either $\pi_1 = \pi_2 = \pi_3$ or else there exist $u \in U$ and $i \in [k]$ such that at least two of $(\pi_1(u))_i = 1$, $(\pi_2(u))_i = 2$, and $(\pi_3(u))_i = 3$ hold.

The motivation for the name "uniquely solvable puzzle" is that a USP can be thought of as a jigsaw puzzle. The puzzle pieces are the sets $\{i : u_i = 1\}$, $\{i : u_i = 2\}$, and $\{i : u_i = 3\}$ with $u \in U$, and the puzzle can be solved by permuting these types of pieces according to π_1, π_2, and π_3, respectively, and reassembling them without overlap into triples consisting of one piece of each of the three types. The definition requires that the puzzle must have a unique solution.

A *strong USP* is a USP in which the defining property is strengthened as follows:

> For all permutations $\pi_1, \pi_2, \pi_3 \in \mathrm{Sym}(U)$, either $\pi_1 = \pi_2 = \pi_3$ or else there exist $u \in U$ and $i \in [k]$ such that *exactly* two of $(\pi_1(u))_i = 1$, $(\pi_2(u))_i = 2$, and $(\pi_3(u))_i = 3$ hold.

One convenient way to depict USPs is by labelling a grid in which the rows correspond to elements of the USP and the columns to coordinates. The ordering of the rows is irrelevant. For example, the following labelling defines a strong USP of size 8 and width 6:

3	3	3	3	3	3
1	3	3	2	3	3
3	1	3	3	2	3
1	1	3	2	2	3
3	3	1	3	3	2
1	3	1	2	3	2
3	1	1	3	2	2
1	1	1	2	2	2

This construction naturally generalizes as follows:

Proposition 3.1. *For each $k \geq 1$, there exists a strong USP of size 2^k and width $2k$.*

Proof. Viewing $\{1,3\}^k \times \{2,3\}^k$ as a subset of $\{1,2,3\}^{2k}$, we define U to be

$$\{u \in \{1,3\}^k \times \{2,3\}^k : \text{for } i \in [k],\, u_i = 1 \text{ iff } u_{i+k} = 2\}.$$

Suppose $\pi_1, \pi_2, \pi_3 \in \mathrm{Sym}(U)$. If $\pi_1 \neq \pi_3$, then there exists $u \in U$ such that $(\pi_1(u))_i = 1$ and $(\pi_3(u))_i = 3$ for some $i \in [k]$. Similarly, if $\pi_2 \neq \pi_3$, then there exists $u \in U$ such that $(\pi_2(u))_i = 2$ and $(\pi_3(u))_i = 3$ for some $i \in [2k] \setminus [k]$. In either case, exactly two of $(\pi_1(u))_i = 1$, $(\pi_2(u))_i = 2$, and $(\pi_3(u))_i = 3$ hold because in each coordinate only two of the three symbols 1, 2, and 3 can occur. It follows that U is a strong USP, as desired. □

We define the *strong USP capacity* to be the largest constant C such that there exist strong USPs of size $(C - o(1))^k$ and width k for infinitely many values of k. (We use the term "capacity" because this quantity is the Sperner capacity of a certain directed hypergraph, as we explain in Section 6.) The *USP capacity* is defined analogously.

There is a simple upper bound for the USP capacity, which is of course an upper bound for the strong USP capacity as well:

Lemma 3.2. *The USP capacity is at most $3/2^{2/3}$.*

Proof. Let U be a USP of width k. For each triple n_1, n_2, n_3 of nonnegative integers summing to k, define the subset U_{n_1,n_2,n_3} of U to consist of all elements of U containing n_1 entries that are 1, n_2 that are 2, and n_3 that are 3. There are $\binom{k+2}{2}$ choices of n_1, n_2, n_3, so

$$|U| \leq \binom{k+2}{2} \max_{n_1,n_2,n_3} |U_{n_1,n_2,n_3}|.$$

If two elements of U have the symbol 1 in exactly the same locations, then letting π_1 interchange them would violate the definition of a USP, and of course the same holds for 2 or 3. Thus,

$$|U_{n_1,n_2,n_3}| \leq \min_i \binom{k}{n_i} \leq \left(\frac{3}{2^{2/3}} + o(1)\right)^k,$$

where the latter inequality holds because $\min_i \binom{k}{n_i}$ is maximized when $n_1 = n_2 = n_3 = k/3$. It follows that $|U| \leq \left(3/2^{2/3} + o(1)\right)^k$, as desired. □

USPs turn out to be implicit in the analysis in Coppersmith and Winograd's paper [3], although they are not discussed as such. Section 6 of [3] can be interpreted as giving a probabilistic construction showing that Lemma 3.2 is sharp:

Theorem 3.3 (Coppersmith and Winograd [3]). *The USP capacity equals $3/2^{2/3}$.*

We conjecture that the same is true for strong USPs:

Conjecture 3.4. *The strong USP capacity equals $3/2^{2/3}$.*

This conjecture would imply that $\omega = 2$, as we explain in the next subsection.

3.2. Using strong USPs

Given a strong USP U of width k, let H be the abelian group of all functions from $U \times [k]$ to the cyclic group Cyc_m (H is a group under pointwise addition). The symmetric group $\mathrm{Sym}(U)$ acts on H via

$$\pi(h)(u,i) = h(\pi^{-1}(u), i)$$

for $\pi \in \mathrm{Sym}(U)$, $h \in H$, $u \in U$, and $i \in [k]$.

Let G be the semidirect product $H \rtimes \mathrm{Sym}(U)$, and define subsets S_1, S_2, and S_3 of G by letting S_i consist of all products $h\pi$ with $\pi \in \mathrm{Sym}(U)$ and $h \in H$ satisfying

$$h(u,j) \neq 0 \qquad \text{iff} \qquad u_j = i$$

for all $u \in U$ and $j \in [k]$.

Proposition 3.5. *If U is a strong USP, then S_1, S_2, and S_3 satisfy the triple product property.*

Proof. Consider a triple product

$$h_1\pi_1\pi_1'^{-1}h_1'^{-1}h_2\pi_2\pi_2'^{-1}h_2'^{-1}h_3\pi_3\pi_3'^{-1}h_3'^{-1} = 1 \qquad (3.1)$$

with $h_i\pi_i, h_i'\pi_i' \in S_i$. For (3.1) to hold we must have

$$\pi_1\pi_1'^{-1}\pi_2\pi_2'^{-1}\pi_3\pi_3'^{-1} = 1. \qquad (3.2)$$

Set $\pi = \pi_1\pi_1'^{-1}$ and $\rho = \pi_1\pi_1'^{-1}\pi_2\pi_2'^{-1}$. Then the remaining condition for (3.1) to hold is that in the abelian group H (with its Sym(U) action),

$$h_1 - h_3' + \pi(h_2 - h_1') + \rho(h_3 - h_2') = 0. \qquad (3.3)$$

Note that

$$\begin{aligned} (h_1 - h_3')(u,j) \neq 0 \quad &\text{iff} \quad u_j \in \{1,3\}, \\ \pi(h_2 - h_1')(u,j) \neq 0 \quad &\text{iff} \quad (\pi^{-1}(u))_j \in \{2,1\}, \text{ and} \\ \rho(h_3 - h_2')(u,j) \neq 0 \quad &\text{iff} \quad (\rho^{-1}(u))_j \in \{3,2\}. \end{aligned}$$

By the definition of a strong USP, either $\pi = \rho = 1$ or else there exist u and j such that exactly one of these three conditions holds, in which case (3.3) cannot hold. Thus, $\pi = \rho = 1$, which together with (3.2) implies $\pi_i = \pi_i'$ for all i. Then we have

$$h_1 + h_2 + h_3 = h_1' + h_2' + h_3',$$

which implies $h_i' = h_i$ for each i (because for different choices of i they have disjoint supports). Thus, the triple product property holds. □

Analyzing this construction using Corollary 1.9 and the bound $[G : H] = |U|!$ on the largest character degree of G yields the following bound:

Corollary 3.6. *If U is a strong USP of width k, and $m \geq 3$ is an integer, then*

$$\omega \leq \frac{3\log m}{\log(m-1)} - \frac{3\log |U|!}{|U|k\log(m-1)}.$$

In particular, if the strong USP capacity is C, then

$$\omega \leq \frac{3(\log m - \log C)}{\log(m-1)}.$$

Proposition 3.1 yields $\omega < 2.67$ with $m = 9$. In the next subsection we prove that the strong USP capacity is at least $2^{2/3}$ and hence $\omega < 2.48$, which is the best bound we know how to prove using strong USPs.

If Conjecture 3.4 holds, then Corollary 3.6 yields $\omega = 2$ upon taking $m = 3$.

3.3. The triangle construction

The strong USP constructed in Proposition 3.1 has the property that only two symbols (of the three possibilities 1, 2, and 3) occur in each coordinate. Every USP with this property is a strong USP, and we can analyze exactly how large such a USP can be as follows.

Suppose $U \subseteq \{1,2,3\}^k$ is a subset with only two symbols occurring in each coordinate. Let H_1 be the subgroup of Sym(U) that preserves the coordinates in which only 1 and 2 occur, H_2 the subgroup preserving the coordinates in which only 2 and 3 occur, and H_3 the subgroup preserving the coordinates in which only 1 and 3 occur.

Lemma 3.7. *The set U is a USP iff H_1, H_2, and H_3 satisfy the triple product property within* Sym(U).

Proof. Suppose $\pi_1, \pi_2, \pi_3 \in$ Sym(U). The permutation $\pi_1\pi_2^{-1}$ is not in H_1 iff there exists $v \in U$ and a coordinate i such that $v_i = 2$ and $((\pi_1\pi_2^{-1})(v))_i = 1$. If we set $u = \pi_2^{-1}(v)$, then this is equivalent to $(\pi_2(u))_i = 2$ and $(\pi_1(u))_i = 1$. Similarly, $\pi_2\pi_3^{-1} \notin H_2$ iff there exist u and i such that $(\pi_2(u))_i = 2$ and $(\pi_3(u))_i = 3$, and $\pi_3\pi_1^{-1} \notin H_3$ iff there exist u and i such that $(\pi_1(u))_i = 1$ and $(\pi_3(u))_i = 3$.

Thus, U is a USP iff for all π_1, π_2, π_3, if $\pi_1\pi_2^{-1} \in H_1$, $\pi_2\pi_3^{-1} \in H_2$, and $\pi_3\pi_1^{-1} \in H_3$, then $\pi_1 = \pi_2 = \pi_3$. That is equivalent to the triple product property for H_1, H_2, and H_3: recall that because these are subgroups, the triple product property says that for $h_i \in H_i$, $h_1h_2h_3 = 1$ iff $h_1 = h_2 = h_3 = 1$. Any three elements h_1, h_2, h_3 satisfying $h_1h_2h_3 = 1$ can be written in the form $h_1 = \pi_1\pi_2^{-1}$, $h_2 = \pi_2\pi_3^{-1}$, and $h_3 = \pi_3\pi_1^{-1}$. □

Proposition 3.8. *For each $k \geq 1$, there exists a strong USP of size $2^{k-1}(2^k+1)$ and width $3k$.*

It follows that the strong USP capacity is at least $2^{2/3}$ and $\omega < 2.48$.

Proof. Consider the triangle

$$\Delta_n = \{(a,b,c) \in \mathbb{Z}^3 : a+b+c = n-1 \text{ and } a,b,c \geq 0\},$$

with $n = 2^k$, and let H_1, H_2, and H_3 be the subgroups of Sym(Δ_n) preserving the first, second, and third coordinates, respectively. By Theorem 1.7, these subgroups satisfy the triple product property in Sym(Δ_n).

To construct the desired strong USP, choose a subset $U \subseteq \{1,2,3\}^{3k}$ as follows. Among the first k coordinates, only 1 and 2 will occur, among the second k only 2 and 3, and among the third k only 1 and 3. In each of these three blocks of k coordinates, there are 2^k possible patterns that be made using the two available symbols. Number these patterns arbitrarily from 0 to $2^k - 1$ (each number will be used for three patterns, one for each pair of symbols). The elements of U will correspond to elements of Δ_n. In particular, the element of U corresponding to $(a,b,c) \in \Delta_n$ will have the a-th pattern in the first k coordinates, the b-th in the second k, and the c-th in the third. It follows from Lemma 3.7 that U is a strong USP. □

One can show using Lemma 3.7 that this construction is optimal:

Corollary 3.9. *If U is a USP of width k such that only two symbols occur in each coordinate, then $|U| \le (2^{2/3} + o(1))^k$.*

The condition of using only two symbols in each coordinate is highly restrictive, but we have been unable to improve on Proposition 3.8. However, we know of no upper bound on the size of a strong USP besides Lemma 3.2, and we see no reason why Conjecture 3.4 should not be true.

4. The simultaneous double product property

There are at least two natural avenues for improving the construction from Subsection 3.3. In the combinatorial direction, one might hope to replace the strong USP of Proposition 3.8 with a larger one; this will reach exponent 2 if Conjecture 3.4 holds. In the algebraic direction, one might hope to keep the combinatorial structure of the triangle construction in place while modifying the underlying group. Such a modification can be carried out using the *simultaneous double product property* defined below, and we conjecture that it reaches $\omega = 2$ as well (Conjecture 4.7).

We say that subsets S_1, S_2 of a group H satisfy the *double product property* if

$$q_1 q_2 = 1 \qquad \text{implies} \qquad q_1 = q_2 = 1,$$

where $q_i \in Q(S_i)$.

Definition 4.1. *We say that n pairs of subsets A_i, B_i (for $1 \le i \le n$) of a group H satisfy the* simultaneous double product property *if*

- *for all i, the pair A_i, B_i satisfies the double product property, and*
- *for all i, j, k,*

  $$a_i (a'_j)^{-1} b_j (b'_k)^{-1} = 1 \qquad \text{implies} \qquad i = k,$$

  where $a_i \in A_i$, $a'_j \in A_j$, $b_j \in B_j$, and $b'_k \in B_k$.

A convenient reformulation is that if one looks at the sets

$$A_i^{-1} B_j = \{a^{-1} b : a \in A_i, b \in B_j\},$$

those with $i = j$ are disjoint from those with $i \ne j$.

For a trivial example, set $H = \mathrm{Cyc}_n^k \times \mathrm{Cyc}_n$, and set $A_i = \{(x, i) : x \in \mathrm{Cyc}_n^k\}$ and $B_i = \{(0, i)\}$. Then the pairs A_i, B_i for $i \in \mathrm{Cyc}_n$ satisfy the simultaneous double product property.

Lemma 4.2. *If n pairs of subsets $A_i, B_i \subseteq H$ satisfy the simultaneous double product property, and n' pairs of subsets $A'_i, B'_i \subseteq H'$ satisfy the simultaneous double product property, then so do the nn' pairs of subsets $A_i \times A'_j, B_i \times B'_j \subseteq H \times H'$.*

Pairs A_i, B_i satisfying the simultaneous double product property in group H can be transformed into subsets satisfying the triple product property via a construction similar to the one in Section 3. Recall that

$$\Delta_n = \{(a,b,c) \in \mathbb{Z}^3 : a + b + c = n - 1 \text{ and } a, b, c \ge 0\}.$$

Given n pairs of subsets A_i, B_i in H for $0 \le i \le n-1$, we define triples of subsets in H^3 indexed by $v = (v_1, v_2, v_3) \in \Delta_n$ as follows:

$$\begin{aligned} \widehat{A}_v &= A_{v_1} \times \{1\} \times B_{v_3} \\ \widehat{B}_v &= B_{v_1} \times A_{v_2} \times \{1\} \\ \widehat{C}_v &= \{1\} \times B_{v_2} \times A_{v_3} \end{aligned}$$

Theorem 4.3. *If n pairs of subsets $A_i, B_i \subseteq H$ (with $0 \le i \le n-1$) satisfy the simultaneous double product property, then the following subsets S_1, S_2, S_3 of $G = (H^3)^{\Delta_n} \rtimes \mathrm{Sym}(\Delta_n)$ satisfy the triple product property:*

$$\begin{aligned} S_1 &= \{\widehat{a}\pi : \pi \in \mathrm{Sym}(\Delta_n), \widehat{a}_v \in \widehat{A}_v \text{ for all } v\} \\ S_2 &= \{\widehat{b}\pi : \pi \in \mathrm{Sym}(\Delta_n), \widehat{b}_v \in \widehat{B}_v \text{ for all } v\} \\ S_3 &= \{\widehat{c}\pi : \pi \in \mathrm{Sym}(\Delta_n), \widehat{c}_v \in \widehat{C}_v \text{ for all } v\} \end{aligned}$$

The proof uses Theorem 1.7 and is similar to the proof of Proposition 3.5; it can be found in the full version of this paper.

Theorem 4.4. *If H is a finite group with character degrees $\{d_k\}$, and n pairs of subsets $A_i, B_i \subseteq H$ satisfy the simultaneous double product property, then*

$$\sum_{i=1}^{n} (|A_i||B_i|)^{\omega/2} \le \left(\sum_k d_k^{\omega}\right)^{3/2}.$$

Using this theorem, the example after Definition 4.1 recovers the trivial bound $\omega \le 3$ as $k \to \infty$.

Proof of Theorem 4.4. Let A'_i, B'_i be the N-fold direct product of the pairs A_i, B_i via Lemma 4.2, and let μ be an arbitrary n-vector of nonnegative integers for which $\sum_{i=1}^{n} \mu_i = N$. Among the pairs A'_i, B'_i are $M = \binom{N}{\mu}$ pairs for which

$$|A'_i||B'_i| = \prod_{i=1}^{n} (|A_i||B_i|)^{\mu_i};$$

call this quantity L. Set $P = |\Delta_M|$, so $P = M(M+1)/2$. The three subsets in Theorem 4.3 each have size $P!L^P$. By Theorem 1.8 and Lemma 1.2 we obtain

$$(P!L^P)^\omega \leq (P!)^{\omega-1}\left(\sum_k d_k^\omega\right)^{3NP}.$$

Taking $2P$-th roots and letting $N \to \infty$ yields

$$\binom{N}{\mu}\left(\prod_{i=1}^{n}(|A_i||B_i|)^{\mu_i}\right)^{\omega/2} \leq \left(\sum_k d_k^\omega\right)^{3N/2}.$$

Finally, we apply Lemma 1.1 with $s_i = (|A_i||B_i|)^{\omega/2}$ and $C = (\sum_k d_k^\omega)^{3/2}$ to obtain the stated inequality. □

It is convenient to use two parameters α and β to describe pairs satisfying the simultaneous double product property: if there are n pairs, choose α and β so that $|A_i||B_i| \geq n^\alpha$ for all i and $|H| = n^\beta$. If H is abelian Theorem 4.4 implies $\omega \leq (3\beta - 2)/\alpha$.

The best construction we know is the following:

Proposition 4.5. *For each $m \geq 2$, there is a construction in $\mathrm{Cyc}_m^{2\ell}$ satisfying the simultaneous double product property with $\alpha = \log_2(m-1) + o(1)$ and $\beta = \log_2 m + o(1)$ as $\ell \to \infty$.*

Taking $m = 6$ yields exactly the same bound as in Subsection 3.3 ($\omega < 2.48$).

Proof. Let $n = \binom{2\ell}{\ell}$. Then $n = 2^{2\ell(1-o(1))}$ so $\beta = \log_2 m + o(1)$. For each subset S of the 2ℓ coordinates of $\mathrm{Cyc}_m^{2\ell}$ with $|S| = \ell$, let A_S be the set of elements that are nonzero in those coordinates and zero in the others. Let $\overline{S}$ denote the complement of S, and set $B_S = A_{\overline{S}}$. For each S, we have $|A_S||B_S| = (m-1)^{2\ell}$, so $\alpha = \log_2(m-1) + o(1)$.

We will show that the pairs A_S, B_S satisfy the simultaneous double product property. Each pair A_S, B_S clearly satisfies the double product property, because the elements of A_S and B_S are supported on disjoint sets of coordinates. Each element of $B_S - A_S$ is nonzero in every coordinate, but if $Q \neq R$ then there is a coordinate in $\overline{R} \cap Q$ (note that this is why we require $|Q| = |R|$). Each element of $B_Q - A_R$ vanishes in that coordinate, so

$$(B_Q - A_R) \cap (B_S - A_S) = \emptyset$$

as desired. □

The only limitations we know of on the possible values of α and β are the following:

Proposition 4.6. *If n pairs of subsets $A_i, B_i \subseteq H$ satisfy the simultaneous double product property, with $|A_i||B_i| \geq n^\alpha$ for all i and $|H| = n^\beta$, then $\alpha \leq \beta$ and $\alpha + 2 \leq 2\beta$.*

Proof. The inequality $\alpha \leq \beta$ follows immediately from the double product property, since that means that the quotient map $(a, b) \mapsto a^{-1}b$ from $A_i \times B_i$ to H is injective.

For the other inequality, first note that $A_1, \dots, A_n$ are disjoint (if $x \in A_i \cap A_j$ with $i \neq j$, and $y \in B_i$, then $x^{-1}y \in (A_i^{-1}B_i) \cap (A_j^{-1}B_i)$, which is impossible). Similarly, $B_1, \dots, B_n$ are also disjoint. It follows that the map

$$\mathrm{Sym}_n \times \mathrm{Sym}_n \times \prod_{i=1}^{n} A_i \times \prod_{i=1}^{n} B_i \to (H^n)^2$$

defined by $(\pi, \rho, a, b) \mapsto (\pi a, \rho b)$ is injective. Here, the group Sym_n acts by permuting the n coordinates. Comparing the sizes of these sets yields $(n!)^2(n^\alpha)^n \leq (n^\beta)^{2n}$, which implies $2\beta \geq \alpha + 2$ as $n \to \infty$. Note that by taking direct powers via Lemma 4.2, one can take n arbitrarily large without changing α and β. □

The most important case is when H is an abelian group. There the bound on ω is $\omega \leq (3\beta - 2)/\alpha$, and Proposition 4.6 shows that the only way to achieve $\omega = 2$ is $\alpha = \beta = 2$. We conjecture that that is possible:

Conjecture 4.7. *For arbitrarily large n, there exists an abelian group H with n pairs of subsets A_i, B_i satisfying the simultaneous double product property such that $|H| = n^{2+o(1)}$ and $|A_i||B_i| \geq n^{2-o(1)}$.*

5. The simultaneous triple product property

Each of our constructions of a group proving a nontrivial bound on ω has the same general form, namely a semidirect product of a permutation group with an abelian group. The crucial part of such a construction is the way in which the abelian part is apportioned among the three subsets satisfying the triple product property.

This apportionment can be viewed as reducing several independent matrix multiplication problems to a single group algebra multiplication, using triples of subsets satisfying the simultaneous triple product property:

Definition 5.1. *We say that n triples of subsets A_i, B_i, C_i (for $1 \leq i \leq n$) of a group H satisfy the* simultaneous triple product property *if*

- *for each i, the three subsets A_i, B_i, C_i satisfy the triple product property, and*
- *for all i, j, k,*

  $$a_i(a'_j)^{-1}b_j(b'_k)^{-1}c_k(c'_i)^{-1} = 1 \quad \text{implies} \quad i = j = k$$

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

We say that such a group simultaneously realizes $\langle |A_1|, |B_1|, |C_1|\rangle, \dots, \langle |A_n|, |B_n|, |C_n|\rangle$.

In most applications the group H will be abelian, in which case it is more conventional to use additive notation. In this notation the implication above becomes

$$a_i - a'_j + b_j - b'_k + c_k - c'_i = 0 \qquad \text{implies} \qquad i = j = k.$$

As an example, let $H = \text{Cyc}_n^3$, and call the three factors H_1, H_2, and H_3. Define

$$A_1 = H_1 \setminus \{0\}, \quad B_1 = H_2 \setminus \{0\}, \quad C_1 = H_3 \setminus \{0\}$$

and

$$A_2 = H_2 \setminus \{0\}, \quad B_2 = H_3 \setminus \{0\}, \quad C_2 = H_1 \setminus \{0\}.$$

This construction is based on the one in Section 2, except that this one is slightly more symmetrical.

Proposition 5.2. *The two triples A_1, B_1, C_1 and A_2, B_2, C_2 satisfy the simultaneous triple product property.*

Proof. Each triple clearly satisfies the triple product property in isolation, so we need only deal with the second condition in the definition. For $i \in \{1, 2\}$ define $U_i = A_i - C_i$, $V_i = B_i - A_i$, and $W_i = C_i - B_i$. What we must prove is that if $u_i + v_j + w_k = 0$ with $u_i \in U_i$, $v_j \in V_j$, and $w_k \in W_k$, then $i = j = k$.

We have

$$U_1 = W_2 = \{(x, 0, z) \in \text{Cyc}_n^3 : x \neq 0, z \neq 0\},$$

$$V_1 = U_2 = \{(x, y, 0) \in \text{Cyc}_n^3 : x \neq 0, y \neq 0\},$$

and

$$W_1 = V_2 = \{(0, y, z) \in \text{Cyc}_n^3 : y \neq 0, z \neq 0\}.$$

If i, j, and k are not all equal, then two of them must be equal but different from the third. In each such case, U_i, V_j, and W_k comprise exactly two of the three subsets of Cyc_n^3 defined in the equations above, with one of those two sets occurring twice. The sum of an element from each cannot vanish, since in the repeated set one coordinate is zero, and the other set is always nonzero in that coordinate. □

The reason for the strange condition in the definition of the simultaneous triple product property is that it is exactly what is needed to reduce several independent matrix multiplications to one group algebra multiplication.

Theorem 5.3. *Let R be any algebra over $\mathbb{C}$. If H simultaneously realizes $\langle n_1, m_1, p_1\rangle, \dots, \langle n_k, m_k, p_k\rangle$, then the number of ring operations required to perform k independent matrix multiplications of sizes $n_1 \times m_1$ by $m_1 \times p_1, \dots,$ $n_k \times m_k$ by $m_k \times p_k$ is at most the number of operations required to multiply two elements of $R[H]$.*

The proof is similar to that of Theorem 1.6:

Proof. Suppose H simultaneously realizes $\langle n_1, m_1, p_1\rangle$, $\dots$, $\langle n_k, m_k, p_k\rangle$ via triples N_i, M_i, P_i with $1 \leq i \leq k$. Let A_i be an $n_i \times m_i$ matrix and B_i an $m_i \times p_i$ matrix. We will index the rows and columns of A_i with the sets N_i and M_i, respectively, those of B_i with M_i and P_i, and those of $A_i B_i$ with N_i and P_i.

Consider the product of the following two elements of $R[H]$:

$$\sum_{i=1}^{k} \sum_{s \in N_i, t \in M_i} (A_i)_{s,t} s^{-1} t$$

and

$$\sum_{i=1}^{k} \sum_{t' \in M_i, u \in P_i} (B_i)_{t',u} t'^{-1} u.$$

We have

$$(s^{-1}t)(t'^{-1}u) = s'^{-1}u'$$

with $s \in N_i$, $t \in M_i$, $t' \in M_j$, $u \in P_j$, $s' \in N_k$, and $u' \in P_k$ iff $i = j = k$ and $s = s'$, $t = t'$, and $u = u'$, so the coefficient of $s^{-1}u$ in the product is

$$\sum_{t \in T} (A_i)_{s,t} (B_i)_{t,u} = (A_i B_i)_{s,u}.$$

Thus, one can simply read off the matrix products from the group algebra product by looking at the coefficients of $s^{-1}u$ with $s \in N_i, u \in P_i$, and the theorem follows. □

Other results about the triple product property also generalize straightforwardly to the simultaneous triple product property, such as the following lemma:

Lemma 5.4. *If n triples of subsets $A_i, B_i, C_i \subseteq H$ satisfy the simultaneous triple product property, and n' triples of subsets $A'_j, B'_j, C'_j \subseteq H'$ satisfy the simultaneous triple product property, then so do the nn' triples of subsets $A_i \times A'_j, B_i \times B'_j, C_i \times C'_j \subseteq H \times H'$.*

By Schönhage's asymptotic sum inequality ((15.11) in [1]), one can deduce a bound on ω from the simultaneous triple product property:

Theorem 5.5. *If a group H simultaneously realizes $\langle a_1, b_1, c_1\rangle, \dots, \langle a_n, b_n, c_n\rangle$ and has character degrees $\{d_k\}$, then*

$$\sum_{i=1}^{n} (a_i b_i c_i)^{\omega/3} \leq \sum_k d_k^{\omega}.$$

Frequently H will be abelian, in which case $\sum_k d_k^{\omega} = |H|$. That occurs in the example from Proposition 5.2, which proves that $\omega < 2.93$ using Theorem 5.5.

In Section 7 we provide a proof of Theorem 5.5 completely within our group-theoretic framework, and show furthermore that any bound on ω that can be achieved using the simultaneous triple product property can also be achieved using the ordinary triple product property. Thus, there is no added generality from the simultaneous triple product property, but it is an important organizing principle.

6. Using the simultaneous triple product property

Every construction we have found of a group proving a nontrivial bound on ω has at its core a simultaneous triple product property construction in an abelian group. Each construction also involves a wreath product, but as explained in Section 7 that is a general tool for dealing with the simultaneous triple product property. Given Theorem 5.5, which can be proved either via the wreath product construction of Section 7 or using the asymptotic sum inequality, one can dispense with non-abelian groups entirely. In this section we explain how to interpret each of our constructions in this setting.

6.1. Local strong USPs

A *local strong USP* of width k is a subset $U \subseteq \{1,2,3\}^k$ such that for each ordered triple $(u,v,w) \in U^3$, with u, v, and w not all equal, there exists $i \in [k]$ such that (u_i, v_i, w_i) is an element of

$$\{(1,2,1),(1,2,2),(1,1,3),(1,3,3),(2,2,3),(3,2,3)\}.$$

Lemma 6.1. *Every local strong USP is a strong USP.*

Proof. Let U be a local strong USP, and suppose $\pi_1, \pi_2, \pi_3 \in \mathrm{Sym}(U)$. If π_1, π_2, and π_3 are not all equal, then there exists $u \in U$ such that $\pi_1(u)$, $\pi_2(u)$, and $\pi_3(u)$ are not all equal. There exists $i \in [k]$ such that $((\pi_1(u))_i, (\pi_2(u))_i, (\pi_3(u))_i)$ is in $\{(1,2,1),(1,2,2),(1,1,3),(1,3,3),(2,2,3),(3,2,3)\}$, and hence exactly two of $(\pi_1(u))_i = 1$, $(\pi_2(u))_i = 2$, and $(\pi_3(u))_i = 3$ hold, as desired. □

The reason for the word "local" is that local strong USPs satisfy a condition for every triple of rows, rather than a weaker global condition on permutations. The advantage of local strong USPs is that they lead naturally to a construction satisfying the simultaneous triple product property:

Theorem 6.2. *Let U be a local strong USP of width k, and for each $u \in U$ define subsets $A_u, B_u, C_u \subseteq \mathrm{Cyc}_\ell^k$ by*

$$\begin{aligned} A_u &= \{x \in \mathrm{Cyc}_\ell^k : x_j \neq 0 \text{ iff } u_j = 1\}, \\ B_u &= \{x \in \mathrm{Cyc}_\ell^k : x_j \neq 0 \text{ iff } u_j = 2\}, \text{ and} \\ C_u &= \{x \in \mathrm{Cyc}_\ell^k : x_j \neq 0 \text{ iff } u_j = 3\}. \end{aligned}$$

Then the triples A_u, B_u, C_u satisfy the simultaneous triple product property.

Note that this construction isolates the key idea behind Proposition 3.5.

Proof. Suppose $u, v, w \in U$ are not all equal and

$$a_u - a'_v + b_v - b'_w + c_w - c'_u = 0$$

with $a_u \in A_u$, $a'_v \in A_v$, $b_v \in B_v$, $b'_w \in B_w$, $c_w \in C_w$ and $c'_u \in C_u$. By the definition of a local strong USP, there exists $i \in [k]$ such that (u_i, v_i, w_i) is in

$$\{(1,2,1),(1,2,2),(1,1,3),(1,3,3),(2,2,3),(3,2,3)\}.$$

In each of these cases exactly one of $a_u, a'_v, b_v, b'_w, c_w, c'_u$ is nonzero, namely a'_v, b_v, a_u, c'_u, b'_w, and c_w, respectively. Thus, in each case the equation $a_u + b_v + c_w = a'_v + b'_w + c'_u$ is impossible, so $u = v = w$, as desired.

All that remains is to show that for each u, the sets A_u, B_u, C_u satisfy the triple product property, which is trivial (they are supported on disjoint sets of coordinates). □

At first glance the definition of a local strong USP appears far stronger than that of a strong USP. For example, the strong USPs constructed in Subsection 3.3 are not local strong USPs. However, it turns out that any bound on ω that can be proved using strong USPs can be proved using local strong USPs:

Proposition 6.3. *The strong USP capacity is achieved by local strong USPs. In particular, given any strong USP U of width k, there exists a local strong USP of size $|U|!$ and width $|U|k$.*

Proof. Let U be a strong USP of width k, and fix an arbitrary ordering $u_1, u_2, \dots, u_{|U|}$ of the elements of U. For each $\pi \in \mathrm{Sym}(U)$, let $U_\pi \in \{1,2,3\}^{|U|k}$ be the concatenation of $\pi(u_1), \pi(u_2), \dots, \pi(u_{|U|})$. Then the set of all vectors U_π is a local strong USP: given any three elements U_{π_1}, U_{π_2}, and U_{π_3} with π_1, π_2, π_3 not all equal, by the definition of a strong USP there exist $u \in U$ and $i \in [k]$ such that exactly two of $(\pi_1(u))_i = 1$, $(\pi_2(u))_i = 2$, and $(\pi_3(u))_i = 3$ hold. Then in the coordinate indexed by u and i, the vectors U_{π_1}, U_{π_2}, and U_{π_3} have entries among $(1,2,1)$, $(1,2,2)$, $(1,1,3)$, $(1,3,3)$, $(2,2,3)$, $(3,2,3)$, as desired. □

Proposition 6.3 explains the choice of the word "capacity": optimizing the size of a local strong USP amounts to determining the Sperner capacity of a certain directed hypergraph (see [8] for background on Sperner capacity). The full version of this paper will explain this perspective more completely.

6.2. Triangle-free sets

The construction in Theorem 4.3 is also easily interpreted in terms of the simultaneous triple product property. Recall the construction of triples $\widehat{A}_v, \widehat{B}_v, \widehat{C}_v$ indexed by $v \in \Delta_n$, defined before Theorem 4.3. These triples almost satisfy the simultaneous triple product property, in the following sense: if

$$a_u(a'_v)^{-1}b_v(b'_w)^{-1}c_w(c'_u)^{-1} = 1$$

then it follows from the simultaneous double product property that $u_1 = w_1$, $v_2 = u_2$, and $w_3 = v_3$. Call a subset S of Δ_n *triangle-free* if for all $u, v, w \in S$ satisfying $u_1 = w_1$, $v_2 = u_2$, and $w_3 = v_3$, it follows that $u = v = w$. Thus, the triples $\widehat{A}_v, \widehat{B}_v, \widehat{C}_v$ with v in a triangle-free subset of Δ_n satisfy the triple product property.

The critical question is whether there is a triangle-free subset of Δ_n of size $|\Delta_n|^{1-o(1)}$. We give a simple construction achieving this using Salem-Spencer sets (see [7]). Let T be a subset of $[\lfloor n/2 \rfloor]$ of size $n^{1-o(1)}$ that contains no three-term arithmetic progression. The following lemma is easily proved:

Lemma 6.4. *The subset $\{(a, b, c) \in \Delta_n : b - a \in T\}$ is triangle-free and has size $|\Delta_n|^{1-o(1)}$.*

6.3. Local USPs and generalizations

USPs also have a local version, just as strong USPs do. A *local USP* is defined analogously to a local strong USP, except that the triple $(1, 2, 3)$ is allowed in addition to $(1, 2, 1)$, $(1, 2, 2)$, $(1, 1, 3)$, $(1, 3, 3)$, $(2, 2, 3)$, and $(3, 2, 3)$. Local USPs are USPs, and they achieve the USP capacity; the proofs are analogous to those for Lemma 6.1 and Proposition 6.3. In what follows we place this construction in a far broader context:

Definition 6.5. *Let H be a finite abelian group. An H-*chart *$\mathcal{C} = (\Gamma, A, B, C)$ consists of a finite set of symbols Γ, together with three mappings $A, B, C : \Gamma \to 2^H$ such that for each $x \in \Gamma$, the sets $A(x), B(x), C(x)$ satisfy the triple product property. Let $\mathcal{H}(\mathcal{C}) \subseteq \Gamma^3$ denote the set of ordered triples (x, y, z) such that*

$$0 \notin A(x) - A(y) + B(y) - B(z) + C(z) - C(x).$$

A local $\mathcal{C}$-USP of width k *is a subset $U \subseteq \Gamma^k$ such that for each ordered triple $(u, v, w) \in U^3$, with u, v, w not all equal, there exists $i \in [k]$ such that $(u_i, v_i, w_i) \in \mathcal{H}(\mathcal{C})$.*

For example, a local USP is a $\mathcal{C}$-USP for the Cyc$_\ell$-chart $\mathcal{C} = (\{1, 2, 3\}, A, B, C)$ with A, B, C defined as follows (below, $\widehat{H} = \text{Cyc}_\ell \setminus \{0, 1\}$):

$$\begin{array}{lll} A(1) = \{0\} & B(1) = -\widehat{H} & C(1) = \{0\} \\ A(2) = \{1\} & B(2) = \{0\} & C(2) = \widehat{H} \\ A(3) = \widehat{H} & B(3) = \{0\} & C(3) = \{0\} \end{array}$$

Theorem 6.6. *Let H be a finite abelian group, $\mathcal{C}$ an H-chart, and U a local $\mathcal{C}$-USP of width k. For each $u \in U$ define subsets $A_u, B_u, C_u \subseteq H^k$ by*

$$A_u = \prod_{i=1}^{k} A(u_i), \quad B_u = \prod_{i=1}^{k} B(u_i), \quad C_u = \prod_{i=1}^{k} C(u_i).$$

Then these triples of subsets satisfy the simultaneous triple product property.

Together with the example above, this theorem gives an analogue of Theorem 6.2 for local USPs. Using Theorem 3.3, this example achieves $\omega < 2.41$.

Using a more complicated chart with 24 symbols, the bound $\omega < 2.376$ from [3] may be derived from Theorem 6.6. For details, see the full version of this paper.

7. The wreath product construction

It remains to prove Theorem 5.5 using purely group-theoretic means. Besides giving a self-contained proof, this will also show that the ordinary triple product property from Definition 1.3 is as strong as the simultaneous triple product property, in the sense that any bound that can be derived from Theorem 5.5 can be proved using Theorem 1.8 as well.

To prove Theorem 5.5, we make use of a wreath product construction. Let H be a group, and define $G = \text{Sym}_n \ltimes H^n$, where the symmetric group Sym_n acts on H^n from the right by permuting the coordinates according to $(h^\pi)_i = h_{\pi(i)}$. We write elements of G as $h\pi$ with $h \in H^n$ and $\pi \in \text{Sym}_n$.

Theorem 7.1. *If n triples of subsets $A_i, B_i, C_i \subseteq H$ satisfy the simultaneous triple product property, then the following subsets H_1, H_2, H_3 of $G = \text{Sym}_n \ltimes H^n$ satisfy the triple product property:*

$$\begin{aligned} H_1 &= \{h\pi : \pi \in \text{Sym}_n, h_i \in A_i \text{ for each } i\} \\ H_2 &= \{h\pi : \pi \in \text{Sym}_n, h_i \in B_i \text{ for each } i\} \\ H_3 &= \{h\pi : \pi \in \text{Sym}_n, h_i \in C_i \text{ for each } i\} \end{aligned}$$

Proof. The proof is analogous to that of Proposition 3.5. Consider a triple product

$$h_1\pi_1\pi_1'^{-1}h_1'^{-1}h_2\pi_2\pi_2'^{-1}h_2'^{-1}h_3\pi_3\pi_3'^{-1}h_3'^{-1} = 1 \qquad (7.1)$$

with $h_i\pi_i, h'_i\pi'_i \in H_i$. (Note that these subscripts index h_1, h_2, h_3, rather than describing coordinates of a single $h \in H$. Once understood that should not cause confusion.) For (7.1) to hold we must have

$$\pi_1\pi_1'^{-1}\pi_2\pi_2'^{-1}\pi_3\pi_3'^{-1} = 1. \qquad (7.2)$$

Set $\pi = \pi_1\pi_1'^{-1}$ and $\rho = \pi_1\pi_1'^{-1}\pi_2\pi_2'^{-1}$. Then the remaining condition for (7.1) to hold is that in the group H^n with its right Sym_n action,

$$h_3'^{-1}h_1\big(h_1'^{-1}h_2\big)^{\pi}\big(h_2'^{-1}h_3\big)^{\rho} = 1$$

In other words, for each coordinate i,

$$\left(h_3'^{-1}\right)_i \left(h_1\right)_i \left(h_1'^{-1}\right)_{\pi(i)} \left(h_2\right)_{\pi(i)} \left(h_2'^{-1}\right)_{\rho(i)} \left(h_3\right)_{\rho(i)} = 1.$$

By the simultaneous triple product property, we find that $\pi(i) = \rho(i) = i$. Thus, $\pi = \rho = 1$, which together with (7.2) implies $\pi_i = \pi_i'$ for all i. Finally, we have

$$h_1 h_1'^{-1} h_2 h_2'^{-1} h_3 h_3'^{-1} = 1,$$

which implies $h_1 = h_1'$, $h_2 = h_2'$, and $h_3 = h_3'$ because each triple A_i, B_i, C_i satisfies the triple product property. □

As a first step towards proving Theorem 5.5, we prove a weaker bound, with the geometric mean replacing the arithmetic mean:

Lemma 7.2. *If H is a finite group with character degrees $\{d_k\}$ and n triples of subsets $A_i, B_i, C_i \subseteq H$ satisfying the simultaneous triple product property, then*

$$n\left(\prod_i (|A_i||B_i||C_i|)^{\omega/3}\right)^{1/n} \le \sum_k d_k^\omega.$$

Proof. The sizes of the three subsets of G in Theorem 7.1 are $n!\prod_i |A_i|$, $n!\prod_i |B_i|$, and $n!\prod_i |C_i|$, respectively. Applying Theorem 1.8 we get the inequality

$$\left((n!)^3 \prod_i |A_i||B_i||C_i|\right)^{\omega/3} \le \sum_j c_j^\omega.$$

By Lemma 1.2 the right-hand-side is at most $(n!)^{\omega-1}\left(\sum_k d_k^\omega\right)^n$, and then dividing both sides by $(n!)^\omega$ yields

$$\left(\prod_i |A_i||B_i||C_i|\right)^{\omega/3} \le (n!)^{-1}\left(\sum_k d_k^\omega\right)^n.$$

This inequality is slightly weaker than the desired inequality, but that is easy to fix by taking direct powers of H via Lemma 5.4. Replacing H with H^t (and n with n^t) yields

$$\left(\prod_i |A_i||B_i||C_i|\right)^{tn^{t-1}\omega/3} \le (n^t!)^{-1}\left(\sum_k d_k^\omega\right)^{tn^t}.$$

Taking tn^t-th roots and letting $t \to \infty$ gives the claimed inequality. □

Proof of Theorem 5.5. Let A_i', B_i', C_i' be the N-fold direct product of the triples A_i, B_i, C_i (via Lemma 5.4), and let μ be an arbitrary n-vector of nonnegative integers for which $\sum_{i=1}^n \mu_i = N$. Among the triples A_i', B_i', C_i' are $\binom{N}{\mu}$ triples for which

$$|A_i'||B_i'||C_i'| = \prod_{i=1}^n (|A_i||B_i||C_i|)^{\mu_i}.$$

Applying Lemma 7.2 to these triples gives

$$\binom{N}{\mu}\prod_{i=1}^n (|A_i||B_i||C_i|)^{\mu_i\omega/3} \le \left(\sum_k d_k^\omega\right)^N.$$

Applying Lemma 1.1 with $s_i = (|A_i||B_i||C_i|)^{\omega/3}$ and $C = \sum_k d_k^\omega$ yields the desired bound. □

Acknowledgements

We are grateful to Michael Aschbacher, Noam Elkies, William Kantor, László Lovász, Amin Shokrollahi, Gábor Simonyi, and David Vogan for helpful discussions.

References

[1] P. Bürgisser, M. Clausen, and M. A. Shokrollahi. *Algebraic Complexity Theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, 1997.

[2] 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`.

[3] D. Coppersmith and S. Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Computation*, 9:251–280, 1990.

[4] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.3*, 2002. `(http://www.gap-system.org)`.

[5] B. Huppert. *Character Theory of Finite Groups*. Number 25 in de Gruyter Expositions in Mathematics. Walter de Gruyter, Berlin, 1998.

[6] G. James and M. Liebeck. *Representations and Characters of Groups*. Cambridge University Press, Cambridge, second edition, 2001.

[7] 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.

[8] G. Simonyi. Perfect graphs and graph entropy. An updated survey. P*erfect graphs*, 293–328, Wiley-Intersci. Ser. Discrete Math. Optim., Wiley, Chichester, 2001.

[9] V. Strassen. Gaussian elimination is not optimal. *Numerische Mathematik*, 13:354–356, 1969.

[10] V. Strassen. Relative bilinear complexity and matrix multiplication. *J. Reine Angew. Math.*, 375/376:406–443, 1987.