Matrix Multiplication II

Yuval Filmus

January 14, 2013

These notes started their life as a lecture given at the Toronto Student Seminar on February 9, 2012. The material is taken mostly from the classic paper by Coppersmith and Winograd [CW]. Other sources are §15.7 of *Algebraic Complexity Theory* [ACT], Stothers's thesis [Sto], V. Williams's recent paper [Wil], and the paper by Cohn at al. [CKSU]. Starred sections are the ones we didn't have time to cover.

We present three different algorithms, all taken from [CW], in rapid succession. All these algorithms are based on Strassen's groundbreaking *laser method*. Strassen's original ideas are described in the appendix.

1 Algorithm 1

The starting point is the identity

$$\begin{aligned}
&\epsilon^3 \sum_{i=1}^{q} \left(x_0^{[0]} y_i^{[1]} z_i^{[1]} + x_i^{[1]} y_0^{[0]} z_i^{[1]} + x_i^{[1]} y_i^{[1]} z_0^{[0]} \right) + O(\epsilon^4) = \\
&\epsilon \sum_{i=1}^{q} (x_0^{[0]} + \epsilon x_i^{[1]})(y_0^{[0]} + \epsilon y_i^{[1]})(z_0^{[0]} + \epsilon z_i^{[1]}) - \\
&\left(x_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} x_i^{[1]} \right) \left(y_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} y_i^{[1]} \right) \left(z_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} z_i^{[1]} \right) + \\
&(1 - q\epsilon) x_0^{[0]} y_0^{[0]} z_0^{[0]}.
\end{aligned}$$

On the left-hand side, we have a sum of three matrix products $\langle 1,1,q\rangle, \langle q,1,1\rangle, \langle 1,q,1\rangle$ sharing some variables. We can indicate these shared variables using superscripts:

$$\epsilon^3(\langle 1,1,q\rangle^{0,1,1} + \langle q,1,1\rangle^{1,0,1} + \langle 1,q,1\rangle^{1,1,0}) + O(\epsilon^4).$$

The notation $\langle n,m,p\rangle^{i,j,k}$ indicates a tensor of type $\langle n,m,p\rangle$ whose x-variables have superscript i, whose y-variables have superscript j, and whose z-variables have superscript k. The point is that different superscripts of x correspond to disjoint sets of variables.

The identity as it stands isn't very useful, since the three (rather trivial) matrix products involved all share variables. The idea of the laser method is to compute the Nth tensor power of the identity, and then somehow separate the variables. The Nth tensor power of the identity has $(q+2)^N$ terms on the right-hand side. What about the left-hand side? It's of the form

$\epsilon^{3N} F + O(\epsilon^{3N+1})$, where F is now the sum of 3^N different matrix product tensors, each of volume q^N (the volume of a tensor $\langle n, m, p \rangle$ is nmp).

The plan now is to take a large subset of the variables so that if we zero out all the other variables, F separates into a disjoint sum. More specifically, we can think of F as a set of triples of indices. We will single out large sets of indices X, Y, Z such that among the triples $F \cap X \times Y \times Z$, no variable repeats. In other words, for each $x \in X$, there is at most one pair $(y, z) \in Y \times Z$ such that $(x, y, z) \in F$.

How many triples can we expect to obtain in this way? Let P be the size of the projection of F into one of the coordinates (they are all the same by symmetry). We definitely cannot expect to get more than P triples. In our case, $P = 2^N$. However, this bound is a tad naive, as the following argument from [CKSU] shows. Classify the triples according to their "source distribution", that is, how many triples of each type $110, 101, 011$ generated them. If there are a, b, c of those, respectively ($a + b + c = N$), then the projections on the coordinates have sizes $\binom{N}{c}, \binom{N}{b}, \binom{N}{a}$. Hence the number of triples of type a, b, c is at most the minimum of these, which is at most $\binom{N}{N/3}$. There are $O(N^2)$ types, hence there can be at most $O(N^2)\binom{N}{N/3} \approx 2^{h(1/3)N}$ triples, where $h(p) = -p \log_2 p - (1-p) \log_2 (1-p)$ is the binary entropy function.

What this argument teaches us is that it is enough to consider one type of triples, in this case $N/3, N/3, N/3$. Technically, this will manifest itself in the following way. For a subset G of F, construct a *conflict graph* by taking the set G as vertices, and connecting any two vertices which share a coordinate. The subset G consisting of all sets of triples of type $N/3, N/3, N/3$ has the property that all vertices have the same degree. Conversely, the degree in the graph corresponding to F depends on the source distribution: a triple of type a, b, c has roughly $2^{N-a} + 2^{N-b} + 2^{N-c}$ neighbors. To see this, consider for example the first coordinate, which consists of c zeroes and $N - c$ ones. Each index equal to one could come from either 110 or 101, while each index equal to zero certainly came from 011.

Surprisingly, this method of reducing F to G by specifying a source distribution is optimal: we can zero variables so that G separates to a disjoint sum of $P^{1-o(1)}$ terms, where P is the projection of G to each of the coordinates. We will state this shortly, but first, a definition.

Definition 1. Let $T \subset \mathbb{Z}^3$ be a finite collection of triples, and let $\phi \colon T \to \mathbb{R}_+$ be arbitrary. We define below the numerical quantity $\mathrm{cap}(T, \phi)$, the *capacity* of T with respect to ϕ.

Let $N \geq 1$. We extend ϕ to T^N multiplicatively, and to subsets of T^N additively. For $i \in \{1, 2, 3\}$, let $\alpha_i \colon T^N \to \mathbb{Z}^N$ denote the projection to the ith coordinate. Say that subsets $X_i \subseteq \alpha_i(T^N)$ ($i \in \{1, 2, 3\}$) are *good* if α_i is injective on $T^N \cap X_1 \times X_2 \times X_3$ for $i \in \{1, 2, 3\}$. The N-capacity $\mathrm{cap}_N(T, \phi)$ is the maximum of $\phi(T^N \cap X_1 \times X_2 \times X_3)$ over all good subsets X_1, X_2, X_3. The capacity $\mathrm{cap}(T, \phi)$ is defined as

$$\mathrm{cap}(T, \phi) = \limsup_{N \to \infty} \sqrt[N]{\mathrm{cap}_N(T, \phi)}.$$

For example, in our case $T = \{(1, 1, 0), (1, 0, 1), (0, 1, 1)\}$, and we can take ϕ to be constant. In the next section, we will see an example in which ϕ has to be non-constant.

We defined the capacity as a limit superior, but a tensor product construction shows that $\mathrm{cap}_{N_1+N_2}(T, \phi) \geq \cap_{N_1}(T, \phi) \cap_{N_2} (T, \phi)$, and this implies that the limit actually exists.

Lemma 1. *Let T, ϕ be as in Definition 1. Suppose that furthermore, each triple in T sums to a constant, that T is closed under permutations, and that ϕ is constant over permutations. Then*

$\mathrm{cap}(T, \phi)$ *satisfies the following inequalities:*

$$\log_2 \mathrm{cap}(T,\phi) \leq \max_{\pi} H(\pi_1) + \sum_{t \in T} \pi(t) \log_2 \phi(t),$$
$$\log_2 \mathrm{cap}(T,\phi) \geq \max_{\pi} H(\pi) - \max_{\tau:\, \tau_1 = \pi_1} H(\tau) + H(\pi_1) + \sum_{t \in T} \pi(t) \log_2 \phi(t),$$

where π ranges over all permutation-invariant probability distributions over T, π_1 is the distribution of the first coordinate under π, and H is the entropy function given by $H(\pi) = -\sum_{t \in T} \pi(t) \log_2 \pi(t)$. Note that the maximum exists since the range of π is compact.

In our present case, T is permutation-invariant, and so the lemma shows that

$$\mathrm{cap}(T, 1) = 2^{H(1/3,2/3)} = 2^{h(1/3)}.$$

We are now in position to apply Schönhage's asymptotic sum inequality. Taking the Nth tensor power of the identity and zeroing out variables appropriately, we get on the left-hand size a disjoint sum of roughly $2^{h(1/3)N}$ terms of volume q^N. Hence

$$2^{h(1/3)N} q^{\tau N} \leq (q+2)^N.$$

Rearranging variables, we deduce

$$\omega \leq 3 \log_q \frac{q+2}{2^{h(1/3)}}.$$

Plugging $q = 8$, we get $\omega \leq 2.404$.

The construction in Lemma 1 is randomized. In other words, the final algorithm isn't explicit. This makes no difference in our model, since it is non-uniform. The method of conditional expectations can determinize the lemma, in case one really cares about uniformity, though *finding* the algorithm, while it takes polynomial time, might take more time than *executing* it.

Tightness of Lemma 1 When the range of T is small, $\tau = \pi$. In particular, this is the case (as can be checked using linear algebra) if all triples in T are non-negative and sum to a constant which is at most 5. For a sum of 6, we have the following counterexample:

$$X = S(0,2,4), S(1,2,3)^2, \quad Y = S(0,3,3)^2, S(1,1,4)^2, S(2,2,2)^6.$$

Here $S(a,b,c)$ is a shorthand for all distinct permutations of (a,b,c). Both X and Y have the same projections into each of the coordinates, but the constituents are different. We have

$$H(X) = H(\tfrac{1}{18}^{[6]}, \tfrac{1}{9}^{[6]}) \approx 3.50, \quad H(Y) = H(\tfrac{1}{9}^{[6]}, \tfrac{1}{3}) \approx 2.64.$$

Coppersmith and Winograd [CW], followed by Williams [Wil], avoid this issue by restricting themselves to distributions π for which $\tau = \pi$. That is, their distributions are specified only by the marginals. As we remarked, in the former case, this doesn't matter. However, in the latter case, it could matter. Indeed, the Nth power of the algorithm contains all non-negative triples (a,b,c) with $a+b+c = 2N$. This property follows from the fact that it holds for the basic algorithm.

1.1 Proof of Lemma 1

Before delving into the proof itself, we need some auxiliary results.

Lemma 2. *Let π be a partition of $n \geq 2$ into k parts. Then there is a constant C_k such that*

$$n^{-C_k} \leq \frac{\binom{n}{\pi}}{2^{H(\pi/n)n}} \leq n^{C_k}.$$

Proof. Without loss of generality, all parts in π are non-empty. According to Wikipedia, there exist constants K_1, K_2 such that

$$K_1 \leq \frac{n!}{n^{n+1/2}e^{-n}} \leq K_2.$$

Therefore

$$\frac{K_1}{K_2^k} \leq \frac{\binom{n}{\pi}}{A} \leq \frac{K_2}{K_1^k},$$

where

$$\begin{aligned} A &= \frac{n^{n+1/2}e^{-n}}{\prod_{p\in\pi} p^{p+1/2}e^{-p}} \\ &= \sqrt{\frac{n}{\prod_{p\in\pi} p}} \prod_{p\in\pi} (n/p)^p \\ &= \sqrt{\frac{n}{\prod_{p\in\pi} p}} 2^{H(\pi/n)n}. \end{aligned}$$

We can estimate

$$\sqrt{n^{1-k}} \leq \frac{A}{2^{H(\pi/n)n}} \leq \sqrt{n}.$$

Therefore

$$\frac{K_1}{K_2^k} n^{(1-k)/2} \leq \frac{\binom{n}{\pi}}{2^{H(\pi/n)n}} \leq \frac{K_2}{K_1^k} n^{1/2}. \qquad \square$$

Lemma 3. *Let π be a finite probability distribution. There exists a constant C such that for all n, there is a finite probability distribution σ such that $n\sigma$ is integral, $|\sigma(t) - \pi(t)| \leq 1/n$ and $|H(\pi) - H(\sigma)| \leq C/n$.*

Proof. Let S be the support of π. Consider σ_L and σ_H defined by $\sigma_L(t) = \lfloor n\pi(t) \rfloor / n$ and $\sigma_H(t) = \lceil n\pi(t) \rceil / n$. Clearly $\sum \sigma_L \leq 1$ while $\sum \sigma_H \geq 1$. It is not hard to conclude that there is a probability distribution σ such that $n\sigma$ is integral and $|\sigma(t) - \pi(t)| \leq 1/n$. Now $\partial H(\pi)/\partial \pi(t) = -\log e(\pi(t)+1)$. In particular, $\|\nabla H(\pi)\|_1 = C$ is constant. It follows that $|H(\pi) - H(\sigma)| \leq C/n$. $\square$

We divide the proof into two parts: upper bound and lower bound.

Upper bound Fix $N \geq 1$, let X_1, X_2, X_3 be good subsets of T^N, and let $F = T^N \cap X_1 \times X_2 \times X_3$. Let D denote the set of all source distributions, that is the set of all functions $d\colon T \to \mathbb{N}$ that sum to N. Note $|D| = \text{poly}(N)$. For each $d \in D$, let $F_d \subset F$ consist of those elements conforming to the distribution d. Also, let d_i denote the projection distribution on the ith component. Since the projections are injective on F, using Lemma 2 we get

$$|F_d| \leq \min_{i=1}^{3} \binom{N}{d_i} \leq \min_{i=1}^{3} 2^{H(d_i/N)N} \,\text{poly}(N).$$

Therefore

$$\begin{aligned}\phi(F) &= \sum_{d \in D} \prod_{t \in T} \phi(t)^{d(t)} |F_d| \\ &\leq \text{poly}(N) \max_{d \in D} \prod_{t \in T} \phi(t)^{d(t)} \min_{i \in \{1,2,3\}} 2^{H(d_i/N)N}.\end{aligned}$$

Now let Δ be the set of all probability distributions on T. For any $d \in D$, $d/N \in \Delta$, and so

$$\text{cap}_N(T, \phi) \leq \text{poly}(N) \max_{\pi \in \Delta} \prod_{t \in T} \phi(t)^{\pi(t)N} \min_{i \in \{1,2,3\}} 2^{H(p_i)N}.$$

Taking Nth roots, we deduce that

$$\text{cap}(T, \phi) \leq \max_{\pi \in \Delta} \prod_{t \in T} \phi(t)^{\pi(t)} \min_{i \in \{1,2,3\}} 2^{H(\pi_i)N}.$$

Given a distribution $\pi \in \Delta$, let $S(\pi)$ denote its symmetrization. Note that $S(\pi)_i = (\pi_1 + \pi_2 + \pi_3)/3$. Since the entropy function is concave,

$$H(S(\pi)_1) \geq \frac{H(\pi_1) + H(\pi_2) + H(\pi_3)}{3} \geq \min(H(\pi_1), H(\pi_2), H(\pi_3)).$$

On the other hand, since ϕ is permutation-invariant,

$$\prod_{t \in T} \phi(t)^{\pi(t)} = \prod_{t \in T} \phi(t)^{S(\pi)(t)}.$$

Therefore the maximum is obtained at some symmetric distribution.

Lower bound This is the difficult part of the proof. We will use Salem-Spencer sets, which are subsets of the integers without three-term arithmetic progressions. Salem and Spencer [SS] showed constructively that for all M, there exist such subsets of $\{1, \ldots, M\}$ of size $M^{1-o(1)}$. Their construction was improved by Behrend [Beh], Moser [Mos] and Elkin [Elk]. Behrend and Elkin's proofs are randomized, while Moser's proof is constructive. For us, the original construction suffices.

In fact, we will employ Salem-Spencer sets over the group $\mathbb{Z}_M$, for M odd. We can construct such sets by taking Salem-Spencer subsets of $\{0, \ldots, (M-1)/2\}$. Such sets still have size $M^{1-o(1)}$.

Without loss of generality, assume that all triples in T sum to zero. Let π be any distribution on T. Fix $N \geq 1$, and let F be the subset of T^N corresponding to source distribution $N\sigma$, where σ is given by Lemma 3.

Let M be an odd prime to be determined later. Let $h\colon \mathbb{Z}_M^n \to \mathbb{Z}_M$ be a random linear function, and choose $x_1, x_2, x_3 \in \mathbb{Z}_M$ randomly under the constraint $x_1 + x_2 + x_3 = 0$. Finally, for $x \in T^N$ define $h_i(x) = h(\alpha_i(x)) + x_i$. Our construction guarantees that if $\alpha_i(x) \neq \alpha_i(y)$ and $i \neq j$ then $(h_i(x), h_i(y), h_j(x))$ is a uniformly random element in $\mathbb{Z}_p^3$ (since M is prime). Finally, define $h_1'(x) = 2h_1(x)$, $h_2'(x) = 2h_2(x)$ and $h_3'(x) = -h_3(x)$; the reason behind this weird definition will become apparent later on. Since M is odd, $(h_i'(x), h_i'(y), h_j'(x))$ is also a uniformly random element.

Let A be some Salem-Spencer set for $\mathbb{Z}_M$ of size $M^{1-o(1)}$. Define sets Y_1, Y_2, Y_3 as follows:

$$Y_i = \{x \in \alpha_i(F) : h_i'(x) \in A\}.$$

By construction, for every $x \in T^N$,

$$h_1(x) + h_2(x) + h_3(x) = h(\alpha_1(x) + \alpha_2(x) + \alpha_3(x)) = 0.$$

Therefore $h_1'(x) + h_2'(x) = 2h_3'(x)$. Since A is a Salem-Spencer set, we deduce that $h_1'(x) = h_2'(x) = h_3'(x)$. Conversely, if $h_1'(x) = h_2'(x)$ then $2h_1'(x) = 2h_3'(x)$ and so $h_1'(x) = h_3'(x)$. This is a property satisfied by every element of $T^n \cap Y_1 \times Y_2 \times Y_3$. This set therefore partitions into sets $(V_a)|_{a \in A}$, where V_a contains those members such that $h_1'(x) = h_2'(x) = h_3'(x) = a$.

We construct the sets X_i out of the set Y_i separately for each a. Given $a \in A$, form a graph whose vertex set is V_a, and two vertices are connected by an edge if they *conflict*, that is, they share one of the three coordinates. Denoting the set of edges by U_a, there are at least $|V_a| - |U_a|$ connected components. Choose a representative from each connected component. The set X_i contains all ith coordinates of representatives. The graph was constructed to guarantee that X_1, X_2, X_3 is good, and the number of triples is $\sum_a |V_a| - |U_a|$.

This construction works for simple examples, in which $F = T^N \cap \alpha_1(F) \times \alpha_2(F) \times \alpha_3(F)$; however, this is not always the case. To refine this construction, let $V_a' = V_a \cap F$, and let $U_a' \subseteq U_a$ consist of those conflicts touching V_a'. When choosing representatives, choose a representative from F is possible. There are at least $|V_a'| - |U_a'|$ connected components containing triple from F (proof: first add the edges in $U_a \setminus U_a'$), and so the resulting number of triples from F is at least $\sum_a |V_a'| - |U_a'|$.

Let us estimate now the sizes of the sets V_a' and U_a'. We need to define a set G related to F:

$$G = \{(x_1, x_2, x_3) : x_i \in \alpha_i(F)\}.$$

We also define $P = |\alpha_1(F)|$. Each element in $x \in F$ belongs to V_a' with probability $1/M^2$, since $h_1'(x), h_2'(x)$ is uniform over $\mathbb{Z}_M^2$, and $h_1'(x) = h_2'(x)$ implies $h_1'(x) = h_2'(x) = h_3'(x)$. Next, we count the number of potential conflicts involving F. For each $x_i \in \alpha_i(F)$, there are $|G|/P$ triples having x_i as their ith coordinate. Hence the number of conflicts involving F is bounded by $3|F||G|/P$. Each conflict (x, y) belongs to U_a' with probability $1/M^3$, since (assuming it is a 1-conflict) $h_1'(x), h_2'(x), h_2'(y)$ is uniform over $\mathbb{Z}_M^3$ (since $\alpha_1(x) = \alpha_1(y)$ and $x \neq y$ implies $\alpha_2(x) \neq \alpha_2(y)$), and so $h_1'(x) = h_2'(x) = h_2'(y) = a$ happens with probability $1/M^3$. Hence the expected value of $|V_a'| - |U_a'|$ is

$$\frac{|F|}{M^2} - \frac{3|F||G|}{PM^3}.$$

Choosing M to be a prime of size roughly $6|G|/P$, the resulting number of triples is at least

$$M^{1-o(1)} \frac{|F|}{2M^2} = M^{-o(1)} \frac{|F|}{2|G|} P.$$

Using Lemma 2 and arguments similar to Lemma 3, we estimate $P \approx 2^{H(\sigma_1)N} \approx 2^{H(\pi_1)N}$. Similarly, $|F| \approx 2^{H(\pi)N}$ and $|G| \approx 2^{H(\tau)N}$, where τ maximizes $H(\tau)$ under the constraint $\tau_i = \pi_i$ for $i \in \{1,2,3\}$; the optimum is permutation-invariant by the concavity of H. Note we estimate $|G|$ by the largest contribution to it since there are only polynomially many source distributions. Putting everything together, we get the statement of the lemma.

Remark An alternative construction, using Strassen's original approach and avoiding Salem-Spencer sets, appears in §15.7 of [ACT]. We did take Strassen's connected components argument, which replaces a less elegant one due to Coppersmith and Winograd.

Open problem Which bound is correct, the lower bound or the upper bound?

2 Algorithm 2

The identity we started from can be slightly tweaked to yield even more:

$$\epsilon^3 \left[\sum_{i=1}^{q} \left(x_0^{[0]} y_i^{[1]} z_i^{[1]} + x_i^{[1]} y_0^{[0]} z_i^{[1]} + x_i^{[1]} y_i^{[1]} z_0^{[0]} \right) + x_0^{[0]} y_0^{[0]} z_{q+1}^{[2]} + x_0^{[0]} y_{q+1}^{[2]} z_0^{[0]} + x_{q+1}^{[2]} y_0^{[0]} z_0^{[0]} \right] + O(\epsilon^4) =$$

$$\epsilon \sum_{i=1}^{q} (x_0^{[0]} + \epsilon x_i^{[1]})(y_0^{[0]} + \epsilon y_i^{[1]})(z_0^{[0]} + \epsilon z_i^{[1]}) -$$

$$\left(x_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} x_i^{[1]} \right) \left(y_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} y_i^{[1]} \right) \left(z_0^{[0]} + \epsilon^2 \sum_{i=1}^{q} z_i^{[1]} \right) +$$

$$(1 - q\epsilon)(x_0^{[0]} + \epsilon^3 x_{q+1}^{[2]})(y_0^{[0]} + \epsilon^3 y_{q+1}^{[2]})(z_0^{[0]} + \epsilon^3 z_{q+1}^{[2]}).$$

The new identity has three more variables $x_{q+1}^{[2]}, y_{q+1}^{[2]}, z_{q+1}^{[2]}$ and three more terms in the left-hand side $\langle 1,1,1 \rangle^{0,0,2} + \langle 1,1,1 \rangle^{0,2,0} + \langle 1,1,1 \rangle^{2,0,0}$. All this at the cost of no new terms in the right-hand side!

We follow the analysis of Algorithm 1, only now we have one more degree of freedom. Suppose p gives weight $\alpha/3$ to each of the terms of volume q, and weight $\beta/3$ to each of the terms of volume 1 (so $\alpha + \beta = 1$). The projection r gives probabilities $(\alpha+2\beta)/3, 2\alpha/3, \beta/3$ to $0, 1, 2$, correspondingly. Following our earlier reasoning, Lemma 1 together with the asymptotic sum inequality yield

$$\omega \leq 3 \log_{q^\alpha} \frac{q+2}{2^{H((\alpha+2\beta)/3, 2\alpha/3, \beta/3)}}.$$

This time the optimization is more difficult, but can be done numerically (or with Lagrange multipliers), giving $q = 6$, $\alpha \approx 0.952$, $\beta \approx 0.048$ and $\omega \leq 2.388$.

Remark The transition from Algorithm 1 to Algorithm 2 involves adding three new variables. In exactly the same way, we could add even more variables, adding for example the term $x_0^{[0]} y_0^{[0]} z_{q+2}^{[3]}$. However, this does not increase the capacity. Indeed, suppose the set X_1 contained two members a, b differing only in z_{q+1} vs. z_{q+2}. Every triple in $X_1 \times X_2 \times X_3$ in which the first coordinate is a has a doppelgänger in which the first coordinate is b, so X_1, X_2, X_3 isn't good.

3 Algorithm 3

Algorithm 3 uses exactly the same identity as Algorithm 2, only squared. There is an added complication, since not all terms are matrix product tensors. We will see how to handle this soon, but first let's see how the identity squared looks like. If we follow the recipe given so far, the indices will be vectors of length 2. Instead, we will take the *sum* of the two components (Coppersmith and Winograd call this *coupling*). This ensures that the x-, y- and z-indices have a constant sum.

3.1 Squared identity

1. The identity squared has $(q+2)^2$ terms on the right-hand side, and 15 terms on the left-hand side, which break up as follows:

(a) 3 terms similar to $\langle 1,1,1\rangle^{0,0,4}$, coming from

$$\langle 1,1,1\rangle^{0,0,2} \otimes \langle 1,1,1\rangle^{0,0,2}.$$

(b) 6 terms similar to $\langle 1,1,2q\rangle^{0,1,3}$, coming from

$$\langle 1,1,q\rangle^{0,1,1} \otimes \langle 1,1,1\rangle^{0,0,2} \oplus \langle 1,1,1\rangle^{0,0,2} \otimes \langle 1,1,q\rangle^{0,1,1}.$$

(c) 3 terms similar to $\langle 1,1,q^2+2\rangle^{0,2,2}$, coming from

$$\langle 1,1,1\rangle^{0,2,0} \otimes \langle 1,1,1\rangle^{0,0,2} \oplus \langle 1,1,1\rangle^{0,0,2} \otimes \langle 1,1,1\rangle^{0,2,0} \oplus \langle 1,1,q\rangle^{0,1,1} \otimes \langle 1,1,q\rangle^{0,1,1}.$$

(d) 3 terms similar to $T_4^{1,1,2}$, coming from

$$\langle 1,q,1\rangle^{1,1,0}\otimes\langle 1,1,1\rangle^{0,0,2}+\langle 1,1,1\rangle^{0,0,2}\otimes\langle 1,q,1\rangle^{1,1,0}+\langle q,1,1\rangle^{1,0,1}\otimes\langle 1,1,q\rangle^{0,1,1}+\langle 1,1,q\rangle^{0,1,1}\otimes\langle q,1,1\rangle^{1,0,1}.$$

The fourth tensor, T_4, is unfortunately not a matrix product tensor in itself, but it can be used to *produce* matrix product tensors, in some sense that we will make precise. We will have to generalize the asymptotic sum inequality from

$$\sum_s (n_s m_s p_s)^\tau \geq \underline{\mathrm{R}}\big(\bigoplus_s \langle n_s, m_s, p_s\rangle\big) \longrightarrow \omega \leq 3\tau \tag{1}$$

to a more general form

$$\sum_s V_\tau(T_s) \geq \underline{\mathrm{R}}\big(\bigoplus_s T_s\big) \longrightarrow \omega \leq 3\tau. \tag{2}$$

3.2 The value of a tensor

We first present the definition of V_τ, and then give some intuition. The proof of (2) is given in the subsection below.

The operator ρ operates on tensors by rotating the variables x, y, z to y, z, x. For example, if

$$T = \langle n,m,p\rangle = \sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{k=1}^{p} x_{ij}y_{jk}z_{ik}$$

then $\rho(T)$, which is obtained by replacing $x \mapsto y$, $y \mapsto z$, $z \mapsto x$, is

$$\rho(T) = \sum_{i=1}^{n}\sum_{j=1}^{m}\sum_{k=1}^{p} x_{ik}y_{ij}z_{jk} = \sum_{k=1}^{p}\sum_{i=1}^{n}\sum_{j=1}^{m} x'_{ki}y_{ij}z'_{kj} = \langle p, n, m\rangle.$$

Here $x'_{ki} = x_{ik}$ and $z'_{kj} = z_{jk}$ denote transposition. Next, we define *symmetrization*:

$$\mathcal{S}(T) = T \otimes \rho(T) \otimes \rho^2(T).$$

So $\mathcal{S}(\langle n, m, p\rangle) = \langle nmp, nmp, nmp\rangle$. Symmetrization is what allowed us to prove $\omega \leq 3\log_{nmp} R(\langle n, m, p\rangle)$.

The Nth tensor power of a tensor T is $T^{\otimes N} = T \otimes \cdots \otimes T$ (N factors). We say that a tensor T *breaks up* into $T' = \bigoplus_s \langle n_s, m_s, p_s\rangle$ if we can zero some variables to get T'. This is the same process we used in Algorithms 1 and 2. We define

$$V_{\tau,N}(T) = \max\left\{\sum_s (n_s m_s p_s)^\tau : T \text{ breaks up into } \bigoplus_s \langle n_s, m_s, p_s\rangle\right\}.$$

For example, $V_{\tau,N}(\langle n, m, p\rangle) = (nmp)^{n\tau}$. Finally,

$$V_\tau(T) = \sup_N V_{\tau,N}(T)^{1/3N}.$$

So $V_\tau(\langle n, m, p\rangle) = (nmp)^\tau$.

The idea behind this definition is that it immediately implies that

$$V_\tau(T) = \underline{\mathrm{R}}(T) \longrightarrow \omega \leq 3\tau. \tag{3}$$

Indeed, suppose that $V_{\tau,N}(T)^{1/3N} \approx \underline{\mathrm{R}}(T)$. So for some n_s, m_s, p_s,

$$V_{\tau,N}(T) = \sum_s (n_s m_s p_s)^\tau,$$

and $\mathcal{S}(T)^{\otimes N}$ reduces to $T' = \bigoplus_s \langle n_s, m_s, p_s\rangle$. The latter implies that

$$\underline{\mathrm{R}}(T') \leq \underline{\mathrm{R}}(\mathcal{S}(T)^{\otimes N}) \leq \underline{\mathrm{R}}(T)^{3N}.$$

The asymptotic sum inequality (1) states that

$$\sum_s (n_s m_s p_s)^\tau \geq \underline{\mathrm{R}}(T') \longrightarrow \omega \leq 3\tau.$$

Since $V_{\tau,N}(T) \approx \underline{\mathrm{R}}(T)^{3N} \geq \underline{\mathrm{R}}(T')$, the condition roughly holds, hence $\omega \leq 3\tau$.

4 Calculating the value

It's time for us to take a closer look at T_4. With the original indices, it looks like this:

$$T_4 = \langle 1, q, 1\rangle^{10,10,02} + \langle 1, q, 1\rangle^{01,01,20} + \langle q, 1, q\rangle^{10,01,11} + \langle q, 1, q\rangle^{01,10,11}.$$

Let's rename the indices to make things clearer:

$$T_4 = \langle 1,q,1\rangle^{3,3,3} + \langle 1,q,1\rangle^{4,4,4} + \langle q,1,q\rangle^{3,4,5} + \langle q,1,q\rangle^{4,3,5}.$$

The tensor T_4 is not symmetric. Before symmetrizing it, we take the Nth power, and only consider terms with $(\alpha/2)N, (\alpha/2)N, (\beta/2)N, (\beta/2)N$ factors of each of the basic types, where $\alpha+\beta=1$. Each such term has volume $q^{(\alpha+2\beta)N}$. The projection into the x-index or the y-index has $N/2$ coordinates equal to 3 and $N/2$ coordinates equal to 4. The projection into the z-index has $(\alpha/2)N$ coordinates equal to 3, $(\alpha/2)N$ coordinates equal to 4, and βN coordinates equal to 5. If we now symmetrize, the volume of each basic term is $q^{3(\alpha+2\beta)N}$, and the projection into each coordinate has size

$$\binom{N}{N/2}^2 \binom{N}{\frac{\alpha}{2}N, \frac{\alpha}{2}N, \beta N}.$$

Using Lemma 1, we get that the value is at least roughly the $3N$th root of

$$\binom{N}{N/2}^2 \binom{N}{\frac{\alpha}{2}N, \frac{\alpha}{2}N, \beta N} q^{3(\alpha+2\beta)N\tau} \approx 2^{2N} 2^{H(\alpha/2,\alpha/2,\beta)N} q^{3(\alpha+2\beta)N\tau}.$$

Taking the Nth root, we get

$$\begin{aligned} V_\tau(T_4)^3 &\geq 4 \cdot 2^{H(\alpha/2,\alpha/2,\beta)} q^{(3\alpha+6\beta)\tau} \\ &= 4(\alpha/2)^{-\alpha}\beta^{-\beta} q^{(3+3\beta)\tau}. \end{aligned}$$

We can find the optimal value of β by substituting $\alpha = 1-\beta$, differentiating the right-hand side with respect to β, and equating to zero. We find that the optimal values of α, β are

$$\alpha = \frac{2}{q^{3\tau}+2}, \beta = \frac{q^{3\tau}}{q^{3\tau}+2}.$$

Substituting this back, we get

$$V_\tau(T_4) \geq 4^{1/3} q^\tau (2+q^{3\tau})^{1/3}.$$

4.1 Analyzing the algorithm

Having calculated the value of T_4, it's time to go back to the main plan. Previously we have optimized the exact fraction to take of each of the basic terms. However we have 15 of these. In the end, we're only really interested in the size of the projection into each coordinate (assuming the projection has the same size for all coordinates). Since there are only 5 values of each coordinate, this reduces the number of degrees of freedom to 4. In fact, since in the identity, the three indices always sum to 4, the average of all components of the x-index should be 4/3. This takes away one more degree of freedom.

Suppose $\alpha+\beta+\gamma+\delta=1$. Take terms in which each term of the first kind appears $(\alpha/3)N$ times, each term of the second type appears $(\beta/6)N$ times, each term of the third type appears $(\gamma/3)N$ times, and each term of the fourth type appears $(\delta/3)N$ times. The normalized histogram of each coordinate will be

$$\frac{2\alpha+\beta+\gamma}{3} : \frac{\beta+2\delta}{3} : \frac{2\gamma+\delta}{3} : \frac{\beta}{3} : \frac{\alpha}{3}.$$

The asymptotic sum inequality equates $(q+2)^2$ with the entropy of this distribution multiplied by the value, which is

$$V_\tau(T_1)^\alpha V_\tau(T_2)^\beta V_\tau(T_3)^\gamma V_\tau(T_4)^\delta.$$

We don't need to explicitly compute the values of T_1, T_2, T_3 since they are matrix product tensors. If we optimize over $q, \alpha, \beta, \gamma, \delta$, we get that $\omega \leq 2.376$.

4.2 Proof*

The generalized form of the asymptotic sum inequality (2) is implicitly used in Coppersmith and Winograd's original paper [CW]. It is stated explicitly in Williams's work [Wil], but not proved there. The treatment below is taken from Stothers's thesis [Sto], who explicitly proves it (we slightly modify his notation).

Lemma 4. *For every two tensors T_1, T_2,*

$$V_{\tau,N}(T_1 \otimes T_2) \geq V_{\tau,N}(T_1) V_{\tau,N}(T_2).$$

As a consequence, for every tensor T,

$$V_{\tau,cN}(T) \geq V_{\tau,N}(T)^c.$$

Proof. Exercise. □

Lemma 5. *For every tensor T,*

$$V_\tau(T) = \lim_{N \to \infty} V_{\tau,N}(T)^{1/3N}.$$

Proof. We defined $V_\tau(T)$ as the supremum of $V_{\tau,N}(T)^{1/3N}$. Pick any $\epsilon > 0$. For some N,

$$V_{\tau,N}(T)^{1/3N} \geq V_\tau(T) - \epsilon.$$

Let $M = cN + r$, where $c \geq 1$ and $0 \leq r < N$. Since

$$\frac{c}{3M} = \frac{1}{3N} - \frac{r}{3MN} > \frac{1}{3N} - \frac{1}{3M},$$

Lemma 4 shows that

$$V_{\tau,M}(T)^{1/3M} \geq V_{\tau,cN}(T)^{1/3M} \geq V_{\tau,N}(T)^{c/3M} \geq (V_\tau(T) - \epsilon) V_\tau(T)^{-1/3M}.$$

If M is large enough,

$$V_{\tau,M}(T)^{1/3M} \geq V_\tau(T) - 2\epsilon.$$ □

Corollary 6. *For every two tensors T_1, T_2,*

$$V_\tau(T_1 \otimes T_2) \geq V_\tau(T_1) V_\tau(T_2).$$

Lemma 7. *For every two tensors T_1, T_2,*

$$V_\tau(T_1 \oplus T_2) \geq V_\tau(T_1) + V_\tau(T_2).$$

Proof. Lemma 5 implies that for large N,

$$V_{\tau,N}(T_1) \approx V_\tau(T_1)^{3N}, \quad V_{\tau,N}(T_2) \approx V_\tau(T_2)^{3N}.$$

The Nth tensor power $\mathcal{S}(T_1 \oplus T_2)^{\otimes N}$ is a disjoint sum of various terms, including

$$\mathcal{S}(T_1)^{N_1} \otimes \mathcal{S}(T_2)^{N_2}$$

for all N_1, N_2 satisfying $N_1 + N_2 = N$. Hence the definition of value implies that

$$V_{\tau,N}(T_1 \oplus T_2) \geq V_{\tau,N_1}(T_1) V_{\tau,N_2}(T_2).$$

If N_1 and N_2 are both large, then

$$V_{\tau,N_1}(T_1) V_{\tau,N_2}(T_2) \approx V_\tau(T_1)^{3N_1} V_\tau(T_2)^{3N_2}.$$

The right-hand side is maximized when $N_1/N_2 \approx V_\tau(T_1)/V_\tau(T_2)$, in which case

$$V_\tau(T_1)^{N_1} V_\tau(T_2)^{N_2} \gtrsim \frac{(V_\tau(T_1) + V_\tau(T_2))^{3N}}{N+1}.$$

When N is large, both of N_1, N_2 are large, and so

$$V_{\tau,N}^{1/3N}(T_1 \oplus T_2) \gtrsim V_\tau(T_1) + V_\tau(T_2). \qquad \square$$

Corollary 8. *The generalized asymptotic sum inequality* (2) *holds.*

5 General formulation*

The estimate of ω obtained using Algorithm 3 is a two-tiered process. The first step is to calculate several *values*, and the second step is to combine these values and get an estimate of ω. Each step requires combining Lemma 1 with the asymptotic sum inequality. We can phrase the method in the following lemma.

Lemma 9. *Let T be a tensor which is a sum of tensors $\{T_i : i \in I\}$, where T_i has indices i. Given a probability distribution p on I, let r_1, r_2, r_3 be the resulting projections on the three coordinates. Then*

$$\log_2 V_\tau(T) \geq \frac{H(r_1) + H(r_2) + H(r_3)}{3} + \mathbb{E}_{i \sim p}[\log_2 V_\tau(T_i)].$$

Proof sketch. Suppose first that the tensors T_i are matrix multiplication tensors. Using the method of proof of Lemma 1, we can demonstrate the existence of subsets $T_N \subset \mathcal{S}(S)^{\otimes N}$ of size very roughly

$$\frac{1}{3N} \log_2 |T_N| \longrightarrow \frac{H(r_1) + H(r_2) + H(r_3)}{3}.$$

Each subset T_N corresponds to a disjoint sum of matrix multiplication tensors $\langle n, m, p \rangle$ such that

$$(nmp)^\tau = \prod_i V_\tau(T_i)^{3p(i)N}.$$

The estimate of $V_\tau(T)$ now follows from the formula for $V_{\tau,N}(T)$.

For the general case, the reader can show that the definition of V_τ remains valid if we replace the original definition of $V_{\tau,N}$ with

$$V_{\tau,N}(T) = \max\left\{\sum_S V_\tau(T_s) : T \text{ breaks up into } \bigoplus_s T_s\right\}. \qquad \square$$

In order to apply the lemma, we need to use the base case stating that $V_\tau(\langle n,m,p\rangle) = (nmp)^\tau$. In order to use the lemma to obtain a bound on ω, we use the asymptotic sum inequality $V_{3\omega}(T) \geq \underline{\mathrm{R}}(T)$.

As an example, consider the tensor T_4 encountered in Algorithm 3. The index set and known values are given by

$$V_\tau(T_{3,3,3}) = V_\tau(T_{4,4,4}) = q^\tau, \quad V_\tau(T_{3,4,5}) = V_\tau(T_{4,3,5}) = q^{2\tau}.$$

Convexity implies that in order to maximize the entropy, we need the probability p to depend only on the value (in this case, only on the volume of the corresponding matrix multiplication tensor), hence we give each of the first two triples probability $\alpha/2$, and each of the latter two triples probability $\beta/2$ (where $\alpha+\beta=1$). We have

$$H(r_1) = H(r_2) = H(1/2,1/2) = 1, \quad H(r_3) = H(\alpha/2,\alpha/2,\beta).$$

Regarding volumes, we have

$$\mathbb{E}_{i\sim p}[\log_2 V_\tau(T_i)] = \tau(\alpha+2\beta)\log_2 q.$$

The lemma now implies

$$\log_2 V_\tau(T_4) \geq \frac{2 + H(\alpha/2,\alpha/2,\beta)}{3} + \tau(\alpha+2\beta)\log_2 q.$$

6 Strassen's version*

Strassen starts with the identity

$$\epsilon\sum_{i=1}^{q}(x_i^{[1]}y_0^{[0]}z_i^{[1]} + x_0^{[0]}y_i^{[1]}z_i^{[1]}) + O(\epsilon^2) =$$
$$\sum_{i=1}^{q}(x_0^{[0]} + \epsilon x_i^{[1]})(y_0^{[0]} + \epsilon y_i^{[1]})z_i - x_0^{[0]}y_0^{[0]}\sum_{i=1}^{q} z_i.$$

On the left, we have $T = \langle q,1,1\rangle^{1,0,1} + \langle 1,1,q\rangle^{0,1,1}$. Strassen in effect calculates $V_\tau(T)$. Raise the identity to the Nth power. The projections have sizes $2^N, 2^N, 1$, so after symmetrizing all the projections have size 4^N. All the individual terms have volume q^{3N}. Lemma 1 (which Strassen proves with an explicit construction) shows that

$$V_\tau(T) \geq 4^{1/3}q^\tau.$$

Therefore, the asymptotic sum inequality shows that $\omega \leq 3\tau$ where

$$4^{1/3}q^{\tau} = q + 1.$$

For $q = 5$, this gives $\omega \leq 2.48$.

One difference between Strassen's account and the later one is the technique used to separate a tensor into disjoint components. Coppersmith and Winograd do this by zeroing out some of the variables. Strassen gives a different ϵ-weight to each input variable, and then obtains an approximate equation (an equation of the type used to define border rank). His method obviates the need for Salem-Spencer sets.

Question: The Coppersmith-Winograd construction can be easily emulated using a Strassen construction. Is the other direction also true?

References

[ACT] Peter Bürgisser, Michael Clausen and M. Amin Shokrollahi, *Algebraic Complexity Theory*, Springer, 1997.

[Beh] Felix Behrend, *On sets of integers which contain no three terms in arithmetic progression*, Proc. Nat. Acad. Sci. 32:331–332, 1946.

[Elk] Michael Elkin, *An improved construction of progression-free sets*, Israeli J. Math. 184:93–128, 2011.

[CKSU] Henry Cohn, Robert Kleinberg, Balázs Szegedy and Chris Umans, *Group-theoretic algorithms for matrix multiplication*, FOCS 2005.

[CW] Dan Coppersmith and Shmuel Winograd, *Matrix multiplication via arithmetic progressions*, J. Symb.Comput. 9(3):251–280, 1990.

[Mos] Leo Moser, *On non-averaging sets of integers*, Canad. J. Math. 5:245–253, 1953.

[SS] Raphaël Salem and Donald Spencer, *On sets of integers which contain no three in arithmetic progressions*, Proc. Nat. Acad. Sci. 28:561–563, 1942.

[Sto] Andrew James Stothers, *On the complexity of matrix multiplication*, Ph.D. thesis (U. of Edinburgh), 2010.

[Wil] Virginia Vassilevska-Williams, *Breaking the Coppersmith-Winograd barrier*, STOC 2012.