Tensor surgery and tensor rank

Matthias Christandl[1] and Jeroen Zuiddam[2]

Abstract

We introduce a method for transforming low-order tensors into higher-order tensors and apply it to tensors defined by graphs and hypergraphs. The transformation proceeds according to a surgery-like procedure that splits vertices, creates and absorbs virtual edges and inserts new vertices and edges. We show that tensor surgery is capable of preserving the low rank structure of an initial tensor decomposition and thus allows to prove nontrivial upper bounds on tensor rank, border rank and asymptotic rank of the final tensors. We illustrate our method with a number of examples. Tensor surgery on the triangle graph, which corresponds to the matrix multiplication tensor, leads to nontrivial rank upper bounds for all odd cycle graphs, which correspond to the tensors of iterated matrix multiplication. In the asymptotic setting we obtain upper bounds in terms of the matrix multiplication exponent ω and the rectangular matrix multiplication parameter α. These bounds are optimal if ω equals two. We also give examples that illustrate that tensor surgery on general graphs might involve the absorption of virtual hyperedges and we provide an example of tensor surgery on a hypergraph. Besides its relevance in algebraic complexity theory, our work has applications in quantum information theory and communication complexity.

1 Introduction

This paper introduces a method for proving upper bounds on tensor rank, border rank and asymptotic tensor rank. The method gives particularly clean results when applied to tensors that are defined combinatorially. Let us first illustrate the combinatorial description that we are using and illustrate the method.

1.1 Illustration

The most famous example of a tensor that fits into our combinatorial framework (and which plays an important role in this paper) is the two-by-two

---

[1]QMATH, Department of Mathematical Sciences, University of Copenhagen, Universitetsparken 5, 2100 Copenhagen, Denmark. Email: christandl@math.ku.dk

[2]QuSoft, CWI Amsterdam and University of Amsterdam, Science Park 123, 1098 XG Amsterdam, Netherlands. Email: j.zuiddam@cwi.nl

matrix multiplication tensor, which is the 3-tensor described by the triangle graph C_3

$$\mathrm{T}_2\Big(\;\triangleright\;\Big) \;=\; \sum_{i \in \{0,1\}^3} (b_{i_1} \otimes b_{i_2}) \otimes (b_{i_2} \otimes b_{i_3}) \otimes (b_{i_3} \otimes b_{i_1}) \;\in\; (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 3},$$

where $\{b_0, b_1\}$ is the standard basis of $\mathbb{C}^2$. For a space $\mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_k}$ of k-tensors we refer to the $\mathbb{C}^{n_i}$ as the tensor legs. Informally, the graph–tensor correspondence is as follows: each vertex of the graph corresponds to a tensor leg and each edge in the graph corresponds to an index to sum over, shared between tensor legs (see Section 2 for a formal definition). By default we view the above tensor as a 3-tensor, but we will sometimes view it as a 6-tensor. Another important example is the so-called rank-two unit 3-tensor, which corresponds to the hypergraph on three vertices with a single hyperedge $\{1,2,3\}$

$$\mathrm{T}_2\Big(\;\triangleright\;\Big) \;=\; \sum_{i \in \{0,1\}} b_i \otimes b_i \otimes b_i \;\in\; \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2.$$

In the algebraic complexity theory literature, the two-by-two matrix multiplication tensor is usually denoted by $\langle 2, 2, 2 \rangle$ and the rank-two unit 3-tensor by $\langle 2 \rangle$. As a final illustrative example consider the complete graph on 4 vertices K_4 and the corresponding 4-tensor

$$\mathrm{T}_2\Big(\;\triangle\;\Big) = \sum_{i \in \{0,1\}^6} (b_{i_1} \otimes b_{i_2} \otimes b_{i_3}) \otimes (b_{i_3} \otimes b_{i_4} \otimes b_{i_5}) \otimes (b_{i_2} \otimes b_{i_4} \otimes b_{i_6})$$
$$\otimes (b_{i_1} \otimes b_{i_5} \otimes b_{i_6}) \;\in\; (\mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 4}.$$

Our aim is to prove nontrivial upper bounds on tensor rank, border rank and asymptotic tensor rank. Let us for now focus on tensor rank. The tensor rank of a k-tensor in $\mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_k}$ is the smallest number r such that the tensor can be written as a sum of r simple tensors $v_1 \otimes \cdots \otimes v_k$ with $v_i \in \mathbb{C}^{n_i}$. The tensor rank of a tensor t is denoted by $\mathrm{R}(t)$. Since tensor rank is invariant under the action of the group $\mathrm{GL}_{n_1} \times \cdots \times \mathrm{GL}_{n_k}$ we will identify tensors that are in the same orbit under this group action.

Going back to our examples, a nontrivial upper bound of seven on the tensor rank of $\langle 2, 2, 2 \rangle$ was obtained by Strassen by constructing an efficient bilinear algorithm for multiplying two-by-two matrices [Str69], a breakthrough result in algebraic complexity theory. The second tensor $\langle 2 \rangle$ is the canonical example of a tensor of rank two. For the third tensor, observe that the graph contains a triangle, and hence Strassen's decomposition of $\langle 2, 2, 2 \rangle$ can directly be upgraded to a nontrivial decomposition of this tensor of size 56. This direct upgrading idea does not work when a tensor corresponds to a graph without triangles, say the five-cycle. The method that we will describe below allows us to prove nontrivial rank upper bounds even for tensors corresponding to graphs that do not contain triangles.

Tensor surgery goes as follows. The central idea is to transform a good decomposition of a well-chosen starting tensor into a good decomposition of a goal tensor. Take a tensor t of which we know (an upper bound on) the tensor rank (or border rank, or asymptotic rank). Then, linearly split up a tensor leg of t into multiple tensor legs and take the tensor product with another tensor s ("inserting s") to obtain our goal tensor, carefully keeping track of the increase in rank that this combined operation causes. Combining our knowledge of the decomposition of t with our knowledge of the rank increase gives a decomposition of the goal tensor.

We illustrate tensor surgery with the 5-tensor of the five-cycle C_5,

$$\mathrm{T}_2\left(\bigpentagon\right) = \sum_{i\in\{0,1\}^5} b_{i_1 i_2} \otimes b_{i_2 i_3} \otimes b_{i_3 i_4} \otimes b_{i_4 i_5} \otimes b_{i_5 i_1} \in (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 5}, \quad (1)$$

where $b_{ij} := b_i \otimes b_j$ with $\{b_0, b_1\}$ the standard basis of $\mathbb{C}^2$. The defining decomposition (1) of $\mathrm{T}_2(C_5)$ has size 32. We can improve this rank upper bound as follows. Define the linear map ϕ by

$$\phi : \mathbb{C}^2 \otimes \mathbb{C}^2 \to (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 3}$$
$$u \otimes v \mapsto \sum_{j\in\{0,1\}^2} (u \otimes b_{j_1}) \otimes (b_{j_1} \otimes b_{j_2}) \otimes (b_{j_2} \otimes v).$$

Let $\psi : (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 3} \to (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 5}$ be the map that applies ϕ at the first tensor leg. Then
$$\mathrm{T}_2(C_5) = \psi(\mathrm{T}_2(C_3)).$$

For $\mathrm{T}_2(C_3)$ we have a good decomposition, namely Strassen's decomposition. Define the elements $b_+ := b_0 + b_1$ and $b_- := b_0 - b_1$ in $\mathbb{C}^2$. For any pair of symbols $x, y \in \{0, 1, +, -\}$ define $b_{xy} := b_x \otimes b_y \in \mathbb{C}^2 \otimes \mathbb{C}^2$. Strassen's decomposition is

$$
\begin{aligned}
\mathrm{T}_2(C_3) \;=\; & - \; b_{-0} \otimes b_{0+} \otimes b_{11} && - \; b_{11} \otimes b_{-0} \otimes b_{0+} && - \; b_{0+} \otimes b_{11} \otimes b_{-0} \\
& + \; b_{-1} \otimes b_{1+} \otimes b_{00} && + \; b_{00} \otimes b_{-1} \otimes b_{1+} && + \; b_{1+} \otimes b_{00} \otimes b_{-1} \\
& + \; (b_{00} + b_{11}) \otimes (b_{00} + b_{11}) \otimes (b_{00} + b_{11}).
\end{aligned}
$$

Applying the linear map ψ to the decomposition yields

$$
\begin{aligned}
\mathrm{T}_2(C_5) &= \psi(\mathrm{T}_2(C_3)) \\
&= - \; \phi(b_{-0}) \otimes b_{0+} \otimes b_{11} - \phi(b_{11}) \otimes b_{-0} \otimes b_{0+} - \phi(b_{0+}) \otimes b_{11} \otimes b_{-0} \\
&\quad + \; \phi(b_{-1}) \otimes b_{1+} \otimes b_{00} + \phi(b_{00}) \otimes b_{-1} \otimes b_{1+} + \phi(b_{1+}) \otimes b_{00} \otimes b_{-1} \\
&\quad + \; \phi(b_{00} + b_{11}) \otimes (b_{00} + b_{11}) \otimes (b_{00} + b_{11}).
\end{aligned}
$$

For any $x, y \in \{0, 1, +, -\}$, $\phi(b_{xy}) = \sum_{j\in\{0,1\}^2} (b_x \otimes b_{j_1}) \otimes (b_{j_1} \otimes b_{j_2}) \otimes (b_{j_2} \otimes b_y)$, which has rank 4 as a 3-tensor. We have $\phi(b_{00} + b_{11}) = \sum_{i\in\{0,1\}^3} (b_{i_1} \otimes b_{i_2}) \otimes$

$(b_{i_2} \otimes b_{i_3}) \otimes (b_{i_3} \otimes b_{i_1})$ for the remaining term, which equals $\mathrm{T}_2(C_3)$ and thus has rank 7 as a 3-tensor, invoking Strassen's decomposition for the second time. Therefore, $\mathrm{R}(\mathrm{T}_2(C_5)) \leq 6 \cdot 4 + 1 \cdot 7 = 31$. This means that we have achieved our goal of constructing a nontrivial rank-31 decomposition of the goal tensor $\mathrm{T}_2(C_5)$, smaller than the trivial decomposition of size 32.

Identifying $\mathrm{T}_2(C_3)$ with the graph C_3, we think of ψ as a "surgery map" that splits a vertex into two vertices and inserts a new vertex together with two edges. In pictures, the effect of applying ψ is

Splitting a tensor leg possibly increases tensor rank, as is the case with the term $b_{00} + b_{11}$ above. To remind us of this, we like to decorate the picture with a "virtual edge" connecting the cut vertices,

The crux is the triangle appearing on the right. This triangle indicates the worst-case situation where $\phi(b_{00} + b_{11}) = \mathrm{T}_2(C_3)$ has rank 7. Of course, to get a good decomposition it is important to also keep track of the best-case situation where $\phi(b_{xy})$ has rank 4.

1.2 Main results

Let $G = (V, E)$ be a graph with vertex set V and edge set E, and let n be a natural number. Let $b_1, \ldots, b_n$ be the standard basis of $\mathbb{C}^n$. We define the order-$|V|$ tensor $\mathrm{T}_n(G)$ as

$$\mathrm{T}_n(G) := \sum_{i \in [n]^E} \bigotimes_{v \in V} \Big(\bigotimes_{\substack{e \in E: \\ v \in e}} b_{i_e} \Big),$$

summing over all tuples i indexed by E with entries in $[n] := \{1, 2, \ldots, n\}$. Let $\mathrm{R}(\mathrm{T}_n(G))$ be the tensor rank of the tensor $\mathrm{T}_n(G)$ and let

$$\omega(\mathrm{T}_2(G)) := \lim_{n \to \infty} \tfrac{1}{n} \log_2 \mathrm{R}(\mathrm{T}_2(G)^{\otimes n}) = \lim_{n \to \infty} \log_n \mathrm{R}(\mathrm{T}_n(G))$$

be the *exponent* of $\mathrm{T}_2(G)$, a measure of the asymptotic behaviour of the tensor rank of $\mathrm{T}_n(G)$. (The limit exists and equals the infimum by Fekete's lemma. The equality follows by relating $\log_{2^n} \mathrm{R}(T_{2^n}(G))$ to $\log_n \mathrm{R}(T_n(G))$.)

The triangle tensor $\mathrm{T}_n(C_3)$ is well-studied, because its tensor rank equals the number of bilinear scalar multiplications required to multiply two $n \times n$ matrices. Moreover the so-called *matrix multiplication exponent*

$$\omega := \omega(\mathrm{T}_2(C_3))$$

4

equals the smallest number $\beta \in \mathbb{R}$ such that for any $\varepsilon > 0$ two $n \times n$ matrices can be multiplied with $\mathcal{O}(n^{\beta+\varepsilon})$ scalar multiplications and additions. A priori, $2 \leq \omega \leq 3$. As mentioned above, Strassen showed that $R(T_2(C_3)) \leq 7$ by constructing an efficient bilinear algorithm for multiplying two-by-two matrices, thereby showing that $\omega \leq \log_2 7$. The rank upper bound was later proven to be tight by Winograd [Win71]. Since Strassen's breakthrough, much effort has been put into obtaining better bounds on ω, the state of the art being $2 \leq \omega < 2.3728639$ [LG14]. Proposed approaches towards obtaining lower bounds on the rank of $T_n(C_3)$ and ω include Strassen's asymptotic spectra [Str91, Str88], the geometric complexity theory programme [Str05, BI11] and Young flattenings [Str83, LO15, LO13]. The recent best upper bounds on ω have been obtained by extending a construction of Coppersmith and Winograd [CW90]. It was shown however that this type of extension cannot prove an upper bound on ω below 2.3078 [AFLG15]. Recently, good upper bounds (not the best) have been obtained by a group-theoretic approach which does not fall under this type of extensions [CU12].

In this paper we go into unexplored terrain by studying the tensor rank and exponent of larger cycle tensors. Our first result is that, for any odd k, the tensor of the k-cycle has a nontrivial tensor rank.[3]

Theorem. *Let k be odd. Then $R(T_2(C_k)) \leq 2^k - 1$.*

This was previously only known for odd $k \leq 5$ [BCZ17]. Let $\omega_k := \omega(T_2(C_k))$. We prove a relationship between the exponents of odd cycles.

Theorem. *Let k, ℓ be odd. Then $\omega_{k+\ell-1} \leq \omega_k + \omega_\ell$.*

We moreover prove an upper bound on the exponent of odd cycles in terms of the dual exponent of matrix multiplication α, which we will define in Section 2.

Theorem. *Let k be odd. Then*

$$\omega_k \leq k - \alpha\Big(1 + \frac{1-\alpha}{k-1+\alpha}\Big) \leq k - \alpha.$$

In particular, since $0.3029805 < \alpha \leq 1$, the exponent ω_k is bounded away from k by a constant.

Our results on the exponent of odd cycles are optimal in the sense that if $\omega = 2$, then $\omega_k = k - 1$ for all odd k. For tensor rank, many open problems remain. As a concrete example, we do not know the value of $R(T_2(C_5))$. We know that it is at least 25 and at most 31 (see Remark 3.2).

Besides looking at graphs, we will in this paper explore tensor surgery on hypergraphs, where one splits up a tensor leg into multiple tensor legs and

---

[3]For even k, trivially $R(T_2(C_k)) = 2^k$, while for odd k, trivially $2^{k-1} \leq R(T_2(C_k)) \leq 2^k$, see Section 2.

instead of a graph inserts a hypergraph. We derive a number of results on the asymptotic rank similar to the ones on cycles graphs.

As the main results indicate, tensor surgery works well for sparse graphs. In a subsequent paper, with an entirely different method, we have obtained nontrivial upper bounds on the exponent of dense graphs [CVZ16] (which in turn can be used again as starting tensors for the tensor surgery put forward in this work). The common theme of that paper and the current paper is the following open problem that generalizes the problem of computing the matrix multiplication exponent ω.

Problem 1.1. Let G be a graph. What is the value of $\omega(\mathrm{T}_2(G))$?

1.3 Connections to other work

Tensor rank has been studied in various fields other than algebraic complexity theory: in algebraic statistics [PS05], in signal processing [CM96], in algebraic geometry in the context of rth secant varieties of Segre varieties [Lan12], in quantum information theory as a monotone for stochastic local operations and classical communication (SLOCC) [CCD+10, VC15], and in communication complexity [DKW11, BCZ17] to characterize the complexity of communication problems, to name a few.

The tensor rank of graph tensors in particular has the following applications. In quantum information language, for any graph G the tensor $\mathrm{T}_n(G)$ is the (unnormalized) quantum state obtained by identifying the vertices of G with quantum systems and letting each edge of the graph correspond to a dimension-n Einstein–Podolsky–Rosen (EPR) pair shared among the vertices contained in the edge. For example, if G contains just a single edge, then $\mathrm{T}_n(G)$ is the EPR pair $\sum_{i=1}^n b_i \otimes b_i$. The notion of a graph tensor naturally extends to hypergraphs. For any hypergraph H consisting of a single edge $\{1, 2, \ldots, k\}$, the tensor $\mathrm{T}_n(H)$ is the (unnormalized) Greenberger–Horne–Zeilinger (GHZ) state $\sum_{i=1}^n |i\rangle^{\otimes k}$ of rank n and order k. Let G be a graph on k vertices. The tensor rank $\mathrm{R}(\mathrm{T}_n(G))$ is the smallest number r such that $\mathrm{T}_r(H)$ can be transformed into $\mathrm{T}_n(G)$ under stochastic local operations and classical communication (SLOCC). The exponent $\omega(\mathrm{T}_n(G))$ is the smallest real number β such that $\lceil \beta + o(n) \rceil$ copies of $\mathrm{T}_2(H)$ can be transformed into n copies of $\mathrm{T}_2(G)$ by SLOCC, when n goes to infinity.

In communication complexity, a notion related to tensor rank called support rank characterizes the so-called nondeterministic quantum communication complexity with quantum broadcast communication of any boolean function [BCZ17]. Graph tensors correspond to the graphwise equality problem, so our upper bounds can be interpreted as upper bounds on the complexity of certain graphwise equality problems. Surprisingly, there is an implication in the other direction, namely upper bounds on the support rank of $\mathrm{T}_n(C_3)$ imply slightly worse upper bounds on the tensor rank of $\mathrm{T}_n(C_3)$.

More precisely, one can define a support rank exponent ω_s analogous to the exponent ω and then the inequality $\omega \leq \frac{3}{2}\omega_s - 1$ holds [CU12]. This tightly connects the study of asymptotic rank to communication complexity and in part motivated the present work. In [BCZ17], an explicit size-31 decomposition of $T_2(C_5)$ was found with computer assistance and nontrivial asymptotic upper bounds were given for all odd k. More precisely, Strassen's laser method combined with the distillation result of [VC17] (see Equation (4) in this paper) was used to get the upper bound $\omega(T_2(C_k)) \leq \min_{q \geq 2} \log_q((q+1)^k/4)$. This bound converges to k when k goes to infinity. The present paper thus answers in the positive the open question of whether ω_k is uniformly bounded away from k.

1.4 Outline of the paper

We will begin by discussing some preliminaries in Section 2. In Section 3 we prove rank upper bounds and exponent upper bounds for odd cycles. In Section 4 we explore the more general hypergraph variant of tensor surgery.

2 Preliminaries

After a formal definition of graph tensors, this section discusses basic notions and results around tensor rank, border rank and asymptotic rank. The sections concludes with a discussion of the lower bounds methods of flattening and Young flattening.

2.1 Graph tensors

All our vector spaces will be complex finite-dimensional vector spaces. However, the ideas in this paper will work over any field. Let $G = (V, E)$ be a graph and let n be a natural number. Let $b_1, \ldots, b_n$ be the standard basis of $\mathbb{C}^n$. We define the $|V|$-tensor $T_n(G)$ as

$$T_n(G) := \sum_{i \in [n]^E} \bigotimes_{v \in V} \Big(\bigotimes_{\substack{e \in E: \\ v \in e}} b_{i_e} \Big),$$

where the sum is over all tuples i indexed by E with entries in the set $[n] := \{1, 2, \ldots, n\}$. Equivalently, we can define $T_n(G)$ as follows:

$$T_n(G) = \bigotimes_{e \in E} \sum_{i \in [n]} (b_i \otimes b_i)_e \otimes (1 \otimes \cdots \otimes 1)_{V \setminus e}.$$

Here the subscripts e and $V \setminus \{e\}$ in a summand indicate that the tensor legs of the summand are permuted by $(1, e_1)(2, e_2)$, and the large tensor product is a tensor product of $|V|$-tensors. We write T for T_2. We can safely ignore the fact that this tensor depends on the choice of order of the edges

and vertices, since the tensor rank does not depend on this order, and we identify tensors that are equivalent up to local GL-action. This definition directly extends to hypergraphs. These tensors were studied in [VC15, VC17] with the notation $\mathrm{GHZ}_n^G = \mathrm{T}_n(G)$. Note that $\mathrm{T}_2(G)^{\otimes k} = \mathrm{T}_2(G^{\cup k}) = \mathrm{T}_{2^k}(G)$ where $G^{\cup k}$ denotes the multigraph obtained from G by taking the union of k copies of G on the same vertex set.

2.2 Tensor rank and exponent

The *tensor rank* of a k-tensor in $\mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_k}$ is the smallest number r such that the tensor can be written as a sum of r simple tensors $v_1 \otimes \cdots \otimes v_k$ with $v_i \in \mathbb{C}^{n_i}$. The tensor rank of a tensor t is denoted by $\mathrm{R}(t)$. When k equals 2, tensor rank is the same as matrix rank and is thus efficiently computable. When k is at least 3, however, deciding tensor rank is NP-hard [Hås90], see also [Shi16] and [SS16] for recent developments. The border rank of t is the smallest number r such that t can be approximated by tensors of rank at most r in the Euclidean topology. We denote border rank by $\underline{\mathrm{R}}(t)$. We refer to [BCS97, Lan12] for an introduction to tensor rank and border rank. We mention in particular that there is an algebraic version of border rank which is also defined over finite fields.

For two k-tensors $s \in U_1 \otimes \cdots \otimes U_k$ and $t \in V_1 \otimes \cdots \otimes V_k$ we say s *restricts to* t, and write $s \geq t$, if there exist linear maps $A_i : U_i \to V_i$ such that $(A_1 \otimes \cdots \otimes A_k)s = t$. Define the asymptotic conversion rate from s to t as

$$\omega(s, t) \coloneqq \lim_{n \to \infty} \frac{1}{n} \min\{m \in \mathbb{N} \mid s^{\otimes m} \geq t^{\otimes n}\}.$$

The minimum of the empty set is defined to be ∞. The limit exists and equals the supremum over n, see Lemma 1.1 in [Str88]. Let $[k]$ denote the hypergraph with vertex set $[k]$ and a single edge containing all vertices. We define the *rank-n unit k-tensor* $\mathrm{T}_n(k)$ as

$$\mathrm{T}_n(k) \coloneqq \mathrm{T}_n([k]) = \sum_{i \in [n]} b_i^{\otimes k}.$$

(So, $\mathrm{T}_2(3) = \langle 2 \rangle$.) The asymptotic log-rank or *exponent* of a tensor t is defined as the limit

$$\omega(t) \coloneqq \omega(\mathrm{T}(k), t) = \lim_{n \to \infty} \frac{1}{n} \min\{m \in \mathbb{N} \mid 2^m \geq \mathrm{R}(t^{\otimes n})\}. \tag{2}$$

The parameter $\omega(t)$ thus measures how many copies of $\mathrm{T}(k)$ are asymptotically needed to create a copy of t by restriction. On the other hand, the parameter $\omega(s, \mathrm{T}(k))^{-1}$ measures how many copies of $\mathrm{T}(k)$ can asymptotically be created from one copy of s by restriction. We call this the subexponent of s.

For any $k \in \mathbb{N}$, let C_k be the cycle graph with vertex set $[k]$ and edge set $\{\{1, 2\}, \{2, 3\}, \ldots, \{k, 1\}\}$. A well-known result is that, asymptotically, two

copies of T(3) can be obtained from the triangle tensor [Str87]:

$$\omega(\mathrm{T}(C_3), \mathrm{T}(3))^{-1} = 2. \tag{3}$$

It was recently shown that this distillation rate holds for all cycles [VC17], that is,

$$\omega(\mathrm{T}(C_k), \mathrm{T}(k))^{-1} = 2 \quad \text{for any } k. \tag{4}$$

See [VC15] and [VC17] for general properties of $\omega(s, t)$. It is an open problem in algebraic complexity theory to compute the exponent of matrix multiplication $\omega = \omega(\langle 2, 2, 2 \rangle) = \omega(\mathrm{T}(C_3))$. The currently best bounds on this number are $2 \leq \omega < 2.3728639$ [LG14]. Rather than improving the bounds on ω, we will in this paper focus on bounding $\omega(\mathrm{T}(C_k))$ for $k > 3$.

We will use the following characterizations of the exponent and subexponent, which are straightforward generalizations of results by Strassen [Str88].

Lemma 2.1. *Let t be a tensor. Then, $\omega(t) = \lim_{N \to \infty} \frac{1}{N} \log_2 \mathrm{R}(t^{\otimes N})$.*

Lemma 2.2. *Let G be a graph on k vertices. Then,*

$$\omega(\mathrm{T}(G)) = \inf\{\beta \in \mathbb{R} \mid \mathrm{R}(\mathrm{T}_n(G)) = \mathcal{O}(n^\beta)\}$$
$$= \inf\{\beta \in \mathbb{R} \mid \mathrm{T}_n(G) \leq \mathrm{T}_{\mathcal{O}(n^\beta)}(k)\}.$$

For any $n \in \mathbb{N}$, $\omega(\mathrm{T}(G)) \leq \log_n \mathrm{R}(T_n(G))$.

Lemma 2.3. *Let G be a graph on k vertices. Then,*

$$\omega(\mathrm{T}(G), \mathrm{T}(k))^{-1} = \sup\{\beta \in \mathbb{R} \mid \mathrm{T}_n(G) \geq \mathrm{T}_{\Omega(n^\beta)}(k)\}.$$

For any $n_1, n_2, n_3 \in \mathbb{N}$, the matrix multiplication tensor $\langle n_1, n_2, n_3 \rangle$ is defined as

$$\langle n_1, n_2, n_3 \rangle := \sum_{i \in [n_1] \times [n_2] \times [n_3]} (b_{i_1} \otimes b_{i_2}) \otimes (b_{i_2} \otimes b_{i_3}) \otimes (b_{i_3} \otimes b_{i_1})$$
$$\in (\mathbb{C}^{n_1} \otimes \mathbb{C}^{n_2}) \otimes (\mathbb{C}^{n_2} \otimes \mathbb{C}^{n_3}) \otimes (\mathbb{C}^{n_3} \otimes \mathbb{C}^{n_1}).$$

So $\langle n, n, n \rangle$ equals $\mathrm{T}_n(C_3)$, and $\langle n_1, n_2, n_3 \rangle$ may be thought of as the tensor corresponding to a triangle with edges "weighted" by n_1, n_2, n_3. For any real numbers $\gamma_1, \gamma_2, \gamma_3 \geq 0$, define

$$\omega(\gamma_1, \gamma_2, \gamma_3) := \inf\{\beta \in \mathbb{R} \mid \mathrm{R}(\langle \lfloor n^{\gamma_1} \rfloor, \lfloor n^{\gamma_2} \rfloor, \lfloor n^{\gamma_3} \rfloor \rangle) = \mathcal{O}(n^\beta)\}. \tag{5}$$

By Lemma 2.2, the definition of $\omega(\gamma_1, \gamma_2, \gamma_3)$ in (5) agrees with the definition of ω in (2) in the sense that $\omega(\langle 2, 2, 2 \rangle) = \omega(1, 1, 1)$. Define the *dual exponent of matrix multiplication* α by

$$\alpha := \sup\{\gamma \in \mathbb{R} \mid \omega(1, 1, \gamma) = 2\}. \tag{6}$$

The currently best bounds on this number are $0.3029805 < \alpha \leq 1$ [LG12]. The number ω equals 2 if and only if α equals 1. The dual exponent α will turn out to be useful in combination with tensor surgery.

We will use the following straightforward property of $\omega(\gamma_1, \gamma_2, \gamma_3)$.

9

Lemma 2.4. *Let $\gamma_1, \gamma_2, \gamma_3, \delta \geq 0$ be real numbers. Then*

$$\omega(\delta\gamma_1, \delta\gamma_2, \delta\gamma_3) = \delta\,\omega(\gamma_1, \gamma_2, \gamma_3).$$

Proof. Suppose $\omega(\gamma_1, \gamma_2, \gamma_3) < \beta$. Then by definition

$$R(\langle \lfloor n^{\gamma_1} \rfloor, \lfloor n^{\gamma_2} \rfloor, \lfloor n^{\gamma_3} \rfloor \rangle) = \mathcal{O}(n^\beta)$$

and thus

$$R(\langle \lfloor n^{\delta\gamma_1} \rfloor, \lfloor n^{\delta\gamma_2} \rfloor, \lfloor n^{\delta\gamma_3} \rfloor \rangle) \leq R(\langle \lfloor N^{\gamma_1} \rfloor, \lfloor N^{\gamma_2} \rfloor, \lfloor N^{\gamma_3} \rfloor \rangle) = \mathcal{O}(n^{\delta\beta})$$

with $n^\delta \leq N \leq n^\delta + 1$ an integer. So $\omega(\delta\gamma_1, \delta\gamma_2, \delta\gamma_3) < \delta\beta$. Conversely, suppose that $\omega(\delta\gamma_1, \delta\gamma_2, \delta\gamma_3) < \beta$. Then by definition

$$R(\langle \lfloor n^{\delta\gamma_1} \rfloor, \lfloor n^{\delta\gamma_2} \rfloor, \lfloor n^{\delta\gamma_3} \rfloor \rangle) = \mathcal{O}(n^\beta)$$

and thus

$$R(\langle \lfloor N^{\gamma_1} \rfloor, \lfloor N^{\gamma_2} \rfloor, \lfloor N^{\gamma_3} \rfloor \rangle) \leq R(\langle \lfloor n^{\delta\gamma_1} \rfloor, \lfloor n^{\delta\gamma_2} \rfloor, \lfloor n^{\delta\gamma_3} \rfloor \rangle) = \mathcal{O}(N^{\beta/\delta}),$$

where n is the smallest integer such that $N \leq n^\delta$. So $\omega(\gamma_1, \gamma_2, \gamma_3) < \beta/\delta$. $\quad\square$

2.3 Lower bound methods

Let $t \in \mathbb{C}^{n_1} \otimes \cdots \otimes \mathbb{C}^{n_k}$. A flattening of t is a grouping of the tensor legs into two groups as to obtain a matrix A_t. The flattening of a simple tensor is a simple matrix (a rank-1 matrix). Therefore, the rank of the flattening matrix A_t is a lower bound for the (border) rank of the tensor t,

$$R(A_t) \leq \underline{R}(t) \leq R(t).$$

Recall that for matrices, rank is multiplicative under taking the tensor product. Therefore, by Lemma 2.1,

$$\log_2 R(A_t) = \omega(A_t) \leq \omega(t).$$

Let $G = (V, E)$ be a graph. A cut of G is a partition of V into two disjoint sets. The size of a cut is the number of edges crossing the cut. A maximum cut is a cut of maximum size. Let $f(G)$ denote the size of a maximum cut of G. Let $V = V_1 \sqcup V_2$ be a cut of G of maximum size $f(G)$. Flattening the tensor $T_n(G)$ along the cut yields the matrix

$$A = \sum_{i \in [n]^E} \left(\bigotimes_{u \in V_1} \left(\bigotimes_{\substack{e \in E: \\ u \in e}} b_{i_e} \right) \right) \otimes \left(\bigotimes_{v \in V_2} \left(\bigotimes_{\substack{e \in E: \\ v \in e}} b_{i_e} \right) \right).$$

The rank of A equals $n^{f(G)}$. Therefore

$$n^{f(G)} = R(A) \leq \underline{R}(T_n(G)) \leq R(T_n(G)), \tag{7}$$

In (7), taking $n = 2$ and taking the logarithm $\log_2$, yields the following inequalities of graph parameters,

$$f(G) \leq \omega(\mathrm{T}(G)) \leq \log_2 \mathrm{R}(\mathrm{T}(G)) \leq |E(G)|. \tag{8}$$

For bipartite graphs, each inequality in (8) is an equality. For odd cycles we get the flattening lower bounds $n^{k-1} \leq \mathrm{R}(\mathrm{T}_n(C_k))$ and $k - 1 \leq \omega(\mathrm{T}_2(C_k))$. There exist more sophisticated flattenings called Young flattenings [LO15], which in our language correspond to a sophisticated splitting of a vertex before flattening. Young flattenings were used in [BCZ17] to show that the flattening lower bound on the border rank of $\mathrm{T}_n(C_k)$ is not tight for odd k. However, we do not know of a Young flattening that improves the asymptotic maximum cut lower bound $f(G) \leq \omega(\mathrm{T}(G))$.

3 Tensor surgery on cycles

In this section we will prove upper bounds on the tensor rank and exponent of cycle tensors using tensor surgery.

3.1 Tensor rank

Let $t = t^1 \otimes \cdots \otimes t^k$ be a simple k-tensor in $\bigotimes_{j=1}^{k}(\mathbb{C}^{a_j} \otimes \mathbb{C}^{b_j})$. Then, for any j, we define the local rank $\mathrm{R}_{\mathbb{C}^{a_j} \otimes \mathbb{C}^{b_j}}(t^j)$ of t^j to be the rank of t^j as an element of $\mathbb{C}^{a_j} \otimes \mathbb{C}^{b_j}$.

Theorem 3.1. *For any odd number $k \geq 3$, the tensor rank of the tensor corresponding to the cycle graph C_k is upper bounded by*

$$\mathrm{R}(\mathrm{T}(C_k)) \leq 2^k - 1.$$

Moreover, $\mathrm{T}(C_k)$ has a decomposition that consists of $2^k - 2$ simple summands whose first tensor leg has local rank 1 and one simple summand whose first tensor leg has local rank 2.

Proof. We prove the statement by induction on odd $k \geq 3$. If $k = 3$, then with notation as in the introduction, Strassen's decompositions is

$$
\begin{aligned}
\mathrm{T}(C_3) \;=\; & - \; b_{-0} \otimes b_{0+} \otimes b_{11} & - \; b_{11} \otimes b_{-0} \otimes b_{0+} & \quad - \; b_{0+} \otimes b_{11} \otimes b_{-0} \\
& + \; b_{-1} \otimes b_{1+} \otimes b_{00} & + \; b_{00} \otimes b_{-1} \otimes b_{1+} & \quad + \; b_{1+} \otimes b_{00} \otimes b_{-1} \\
& + \; (b_{00} + b_{11}) \otimes (b_{00} + b_{11}) \otimes (b_{00} + b_{11}),
\end{aligned}
$$

so the statement of the theorem holds. Assume that the statement holds for $k = \ell$. This means that $\mathrm{T}(C_\ell) = \sum_{i=1}^{2^\ell - 1} t_i^1 \otimes \cdots \otimes t_i^\ell$ for some $t_i^j \in \mathbb{C}^2 \otimes \mathbb{C}^2$ such that

$$\#\{i \mid \mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 1\} = 2^\ell - 2 \quad \text{and} \quad \#\{i \mid \mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 2\} = 1.$$

11

Define the linear map ϕ by

$$\phi : \mathbb{C}^2 \otimes \mathbb{C}^2 \to (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes 3}$$

$$u \otimes v \mapsto \sum_{j \in \{0,1\}^2} (u \otimes b_{j_1}) \otimes (b_{j_1} \otimes b_{j_2}) \otimes (b_{j_2} \otimes v).$$

Let $\psi_\ell : (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes \ell} \to (\mathbb{C}^2 \otimes \mathbb{C}^2)^{\otimes \ell + 2}$ be the map that applies ϕ at the first tensor leg. Then

$$\mathrm{T}(C_{\ell+2}) = \psi_\ell(\mathrm{T}(C_\ell)) = \sum_{i=1}^{2^\ell - 1} \phi(t_i^1) \otimes t_i^2 \otimes \cdots \otimes t_i^\ell.$$

If $\mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 1$, then $\phi(t_i^1)$ has a decomposition of size 4 such that for every simple summand the first tensor leg has local rank 1. If $\mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 2$, then $\phi(t_i^1) \cong \mathrm{T}(C_3)$ has a decomposition of size 7 such that for six simple summands the first tensor leg has local rank 1, while for one simple summand the first tensor leg has local rank 2. We conclude that $\mathrm{T}(C_{\ell+2})$ has rank at most $(2^\ell - 2)4 + 1 \cdot 7 = 2^{\ell+2} - 1$. Moreover, $\mathrm{T}(C_{\ell+2})$ has a decomposition that consists of $2^{\ell+2} - 2$ simple summands whose first tensor leg has local rank 1 and one simple summand whose first tensor leg has local rank 2. We conclude that the statement of the theorem holds for $k = \ell + 2$. $\qquad \square$

Remark 3.2. Before moving on, let us say something about lower bounds on the tensor rank $\mathrm{R}(\mathrm{T}(C_k))$. For any k, instead of flattening $\mathrm{T}(C_k)$ to a matrix, we can flatten $\mathrm{T}(C_k)$ to the 3-tensor $\langle 2, 2, 2^{k-2} \rangle$. As mentioned in [BCZ17], by taking a Young flattening of the latter, we get the lower bound

$$2^k - 2^{k-2} + 1 \le \underline{\mathrm{R}}(\mathrm{T}(C_k)).$$

Applying the rank lower bound $\mathrm{R}(\langle n, n, m \rangle) \ge 2mn + 2n - m - 2$ for $m \ge n \ge 3$ of [Blä03] to $\langle 2, 2, 2^{k-2} \rangle$ gives

$$2^k - 2^{k-2} + 2 \le \mathrm{R}(\mathrm{T}(C_k)). \tag{9}$$

For the triangle graph, Strassen already showed that $\mathrm{R}(\mathrm{T}(C_3)) \le 7$ [Str69] and Winograd showed that $\mathrm{R}(\mathrm{T}(C_3)) \ge 7$ [Win71]. Only quite recently Landsberg showed that also $\underline{\mathrm{R}}(\mathrm{T}(C_3)) \ge 7$ [Lan06]. For the next smallest interesting case $\mathrm{T}(C_5)$, Theorem 3.1 brings the rank and border rank in the following ranges:

$$24 \le \underline{\mathrm{R}}(\mathrm{T}(C_5)) \le 31,$$
$$25 \le \mathrm{R}(\mathrm{T}(C_5)) \le 31.$$

Remark 3.3. We mention that the decomposition of $\mathrm{T}(C_5)$ given by the proof of Theorem 3.1 is different from the decomposition given in [BCZ17] in the sense of De Groote's work [dG78], that is, the decompositions can not be transformed into each other by sandwiching and cyclic permutation. This is because the local ranks of the summands are incompatible.

In general, with tensor surgery we can transform decompositions of matrix multiplication tensors $\langle n_1, n_2, n_3 \rangle$ to decompositions of $\mathrm{T}_n(C_k)$. We will illustrate this with $\mathrm{T}_4(C_5)$. We make use of the bounds $\mathrm{R}(\langle 4,4,2 \rangle) \leq 26$, $\mathrm{R}(\langle 4,4,4 \rangle) \leq 49$ and $\underline{\mathrm{R}}(\langle 4,4,2 \rangle) \leq 24$, $\underline{\mathrm{R}}(\langle 4,4,4 \rangle) \leq 46$, see [HK71, Smi13]. First note that $\mathrm{R}(\mathrm{T}_4(C_5)) \leq \mathrm{R}(\mathrm{T}_2(C_5))^2 \leq 31^2$.

Proposition 3.4. $\mathrm{R}(\mathrm{T}_4(C_5)) \leq 937 < 31^2$ *and* $\underline{\mathrm{R}}(\mathrm{T}_4(C_5)) \leq 910$.

Proof. Let ϕ be the linear map $\mathbb{C}^4 \otimes \mathbb{C}^4 \mapsto (\mathbb{C}^4 \otimes \mathbb{C}^4)^{\otimes 3}$ defined on simple tensors by $u \otimes v \mapsto \sum_{j \in [4]^2} (u \otimes b_{j_1}) \otimes (b_{j_1} \otimes b_{j_2}) \otimes (b_{j_2} \otimes v)$, and let ψ be the linear map $(\mathbb{C}^4 \otimes \mathbb{C}^4)^{\otimes 3} \to (\mathbb{C}^4 \otimes \mathbb{C}^4)^{\otimes 5}$ which applies ϕ to the first tensor leg. Then $\mathrm{T}_4(C_5)$ equals $\psi(\langle 4,4,4 \rangle)$. Taking the tensor square of Strassen's decomposition gives a decomposition $\sum_{i=1}^{49} t_i^1 \otimes \cdots \otimes t_i^k$ of $\langle 4,4,4 \rangle$ such that

$$\#\{i \mid \mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 1\} = 6^2,$$
$$\#\{i \mid \mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 2\} = 6 + 6,$$
$$\#\{i \mid \mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 4\} = 1.$$

If $\mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 1$, then $\phi(t_i^1)$ has rank 4^2. If $\mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 2$, then $\phi(t_i^1) \cong \langle 4,4,2 \rangle$ has rank at most 26. If $\mathrm{R}_{\mathbb{C}^2 \otimes \mathbb{C}^2}(t_i^1) = 4$, then $\phi(t_i^1) \cong \langle 4,4,4 \rangle$ has rank at most 49. Therefore, applying ψ to the simple summands $t_i^1 \otimes \cdots \otimes t_i^k$ we obtain $\mathrm{R}(\mathrm{T}_4(C_5)) \leq 6^2 \cdot 4^2 + 12 \cdot 26 + 1 \cdot 49 = 937$.

For the border rank, we have $\underline{\mathrm{R}}(\langle 4,4,2 \rangle) \leq 24$ and $\underline{\mathrm{R}}(\langle 4,4,4 \rangle) \leq 46$, so that by the same argument $\underline{\mathrm{R}}(\mathrm{T}_4(C_5)) \leq 6^2 \cdot 4^2 + 12 \cdot 24 + 1 \cdot 46 = 910$. $\square$

3.2 Exponent

In view of the lower bound (9), our tensor rank bounds might not look that strong. We will now see, however, that applying the same techniques in the asymptotic setting yields optimal bounds, assuming $\omega = 2$. Let $\omega_k \coloneqq \omega(\mathrm{T}(C_k))$.

Theorem 3.5. *For k, ℓ odd, $\omega_{k+\ell-1} \leq \omega_k + \omega_\ell$.*

The idea of the proof is to take the k-cycle C_k, split one vertex in C_k into two vertices and insert $\ell - 2$ new vertices in the graph together with the appropriate $\ell - 1$ edges in order to create the $(k + \ell - 1)$-cycle. In pictures, for $k = 5$ and $\ell = 3$,

Next we consider an optimal decomposition of $\mathrm{T}_n(C_k)$. Not only inserting $\ell - 1$ edges comes with a cost, but also splitting the vertex. The crucial observation is that the total cost is at most the cost of creating an ℓ-cycle, which is asymptotically ω_ℓ.

Proof. Let ϕ be the linear map $\mathbb{C}^n \otimes \mathbb{C}^n \to (\mathbb{C}^n \otimes \mathbb{C}^n)^{\otimes \ell}$ defined on simple tensors by $u \otimes v \mapsto \sum_{j \in [n]^{\ell-1}} (u \otimes b_{j_1}) \otimes (b_{j_1} \otimes b_{j_2}) \otimes \cdots \otimes (b_{j_{\ell-1}} \otimes v)$, and let ψ be the linear map $(\mathbb{C}^n \otimes \mathbb{C}^n)^{\otimes k} \to (\mathbb{C}^n \otimes \mathbb{C}^n)^{\otimes k+\ell-1}$ that applies ϕ at the first tensor leg. Then $\mathrm{T}_n(C_{k+\ell-1}) = \psi(\mathrm{T}_n(C_k))$. Let $\varepsilon > 0$. Then there is a constant $c_\varepsilon \in \mathbb{N}$ and a decomposition of $\mathrm{T}_n(C_k)$ as a sum of at most $c_\varepsilon n^{\omega_k + \varepsilon}$ simple summands (Lemma 2.2). Consider one simple summand $t^1 \otimes \cdots \otimes t^k$ in this decomposition. We have $\mathrm{R}_{\mathbb{C}^n \otimes \mathbb{C}^n}(t^1) \leq n$ and hence $\phi(t^1) \leq \mathrm{T}_n(C_\ell)$. The rank of $\psi(t^1 \otimes \cdots \otimes t^k)$ is therefore at most $d_\varepsilon n^{\omega_\ell + \varepsilon}$ for some constant $d_\varepsilon \in \mathbb{N}$. We conclude that the rank of $\psi(\mathrm{T}_n(C_k))$ is at most $c_\varepsilon d_\varepsilon n^{\omega_k + \omega_\ell + 2\varepsilon}$, and thus $\omega_{k+\ell-1} \leq \omega_k + \omega_\ell$ (Lemma 2.2). $\qquad\square$

Corollary 3.6. *Let $k \geq 5$ odd. Then, $\omega_k \leq \omega_{k-2} + \omega_3$ and thus $\omega_k \leq \frac{k-1}{2}\omega$.*

Corollary 3.7. *If $\omega = 2$, then $\omega_k = k - 1$ for all odd k.*

Remark 3.8. The proofs of Theorem 3.1 and Proposition 3.4 crucially relied on a careful local rank analysis of Strassen's decomposition and other decompositions of matrix multiplication tensors. The same technique may be applied in the asymptotic setting to improve the results of Theorem 3.5, in the following sense. Suppose one has a specific upper bound for ω_k together with information about the local ranks in the corresponding decomposition of $\mathrm{T}_n(C_k)$ for any n. Then, when applying the surgery map ψ to such a decomposition, as in the proof of Theorem 3.5, one can use the specific local rank information instead of using the worst-case upper bound $\mathrm{R}_{\mathbb{C}^n \otimes \mathbb{C}^n}(t^1) \leq n$, and thus obtain an improved asymptotic bound.

The local rank viewpoint reveals an interesting fact about the decompositions of cycle tensors, which is also relevant for the asymptotic local rank analysis idea. Namely, take the tensor $\mathrm{T}_n(C_k)$ and let ψ be the map that split one of the vertices,

$$\psi : (\mathbb{C}^n \otimes \mathbb{C}^n)^{\otimes k} \to (\mathbb{C}^n \otimes \mathbb{C}^n)^{\otimes k-1} \otimes \mathbb{C}^n \otimes \mathbb{C}^n.$$

Then $\psi(\mathrm{T}_n(C_k)) = \mathrm{T}_n(L_k)$ where L_k is the linear graph with k edges, and hence we have $\mathrm{R}(\psi(\mathrm{T}_n(C_k))) = n^k$. Therefore, if $\mathrm{T}_n(C_k) = \sum_{i=1}^r t_i^1 \otimes \cdots \otimes t_i^k$ is a decomposition into simple tensors, then for any $j \in [k]$ we have

$$\sum_{i=1}^r \mathrm{R}_{\mathbb{C}^n \otimes \mathbb{C}^n}(t_i^j) \geq n^k.$$

Let $r = n^\beta$ and let $\mathrm{T}_n(C_k) = \sum_{i=1}^r t_i^1 \otimes \cdots \otimes t_i^k$ be a decomposition. Then the average local rank at the jth leg is lower bounded by

$$\frac{1}{r} \sum_{i=1}^r \mathrm{R}_{\mathbb{C}^n \otimes \mathbb{C}^n}(t_i^j) \geq n^{k-\beta},$$

while $\max_{i\in[r]} \mathrm{R}_{\mathbb{C}^n\otimes\mathbb{C}^n}(t_i^j) \leq n$. If β is close to $k-1$, then the average local rank is close to the maximum. However, if β is bounded away from $k-1$ then there is a gap between average and maximum local rank, so that an improvement by local rank analysis as described above is possible.

The following theorem gives an upper bound on ω_k in terms of the dual exponent of matrix multiplication α.

Theorem 3.9. *For any odd $k \geq 3$, $\omega_k \leq k - \alpha\big(1 + \frac{1-\alpha}{k-1+\alpha}\big) \leq k - \alpha$.*

The idea of the proof is as follows. Start with the unbalanced triangle tensor $\langle n, n, \lfloor n^\alpha \rfloor \rangle$. On the graph level, we split a vertex, and insert a vertex with two edges:

The crucial observation is that the total cost of splitting a vertex and inserting one vertex with the two appropriate edges is $\omega(\langle n, n, \lfloor n^\alpha \rfloor \rangle)$ which is 2. Repeating this procedure $(k-1)/2$ times yields $\mathrm{T}_n(C_k)$ but with edges "weighted" by $\lfloor n^\alpha \rfloor, n, \dots, n$ respectively, at cost $k-1$ in the exponent. To get an evenly weighted $\mathrm{T}_n(C_k)$ we symmetrise cyclically.

Proof. Let $0 < \gamma < \alpha$. Let $\mathrm{T}_{n,\gamma}(C_\ell)$ be the cycle tensor with edges weighted by $\lfloor n^\gamma \rfloor, n, \dots, n$ respectively,

$$\mathrm{T}_{n,\gamma}(C_\ell) = \sum_{i\in[\lfloor n^\gamma \rfloor]\times[n]^{\times(\ell-1)}} (b_{i_1}\otimes b_{i_2})\otimes(b_{i_2}\otimes b_{i_3})\otimes\cdots\otimes(b_{i_\ell}\otimes b_{i_1}).$$

We will show that $\mathrm{R}(\mathrm{T}_{n,\gamma}(C_k)) = \mathcal{O}(n^{k-1+\varepsilon})$ for all $\varepsilon > 0$ by induction on odd $k \geq 3$. For $k = 3$, the statement is true by definition of α. Suppose the statement holds for $k = \ell$. Let ϕ be the linear map $\mathbb{C}^{\lfloor n^\gamma \rfloor}\otimes\mathbb{C}^n \to (\mathbb{C}^{\lfloor n^\gamma \rfloor}\otimes\mathbb{C}^n)\otimes(\mathbb{C}^n\otimes\mathbb{C}^n)^{\otimes 2}$ defined on simple tensors by

$$u\otimes v \mapsto \sum_{j\in[n]^2} (u\otimes b_{j_1})\otimes(b_{j_1}\otimes b_{j_2})\otimes(b_{j_2}\otimes v),$$

and let ψ_ℓ be the linear map $(\mathbb{C}^{\lfloor n^\gamma \rfloor}\otimes\mathbb{C}^n)\otimes(\mathbb{C}^n\otimes\mathbb{C}^n)^{\otimes\ell-2}\otimes(\mathbb{C}^n\otimes\mathbb{C}^{\lfloor n^\gamma \rfloor}) \to (\mathbb{C}^{\lfloor n^\gamma \rfloor}\otimes\mathbb{C}^n)\otimes(\mathbb{C}^n\otimes\mathbb{C}^n)^{\otimes\ell}\otimes(\mathbb{C}^n\otimes\mathbb{C}^{\lfloor n^\gamma \rfloor})$ that applies ϕ at the first tensor leg. Then,

$$\mathrm{T}_{n,\gamma}(C_{\ell+2}) = \psi_\ell(\mathrm{T}_{n,\gamma}(C_\ell)).$$

Let $\varepsilon > 0$. There is a constant $c_\varepsilon \in \mathbb{N}$ and a decomposition of $\mathrm{T}_{n,\gamma}(C_\ell)$ as a sum of at most $c_\varepsilon n^{(\ell-1)+\varepsilon}$ simple summands (Lemma 2.2). Consider one simple summand $t^1\otimes\cdots\otimes t^\ell$ in this decomposition. We have $\mathrm{R}_{\mathbb{C}^{\lfloor n^\gamma \rfloor}\otimes\mathbb{C}^n}(t^1) \leq \lfloor n^\gamma \rfloor$ and hence $\phi(t^1) \leq \mathrm{T}_{n,\gamma}(C_3)$. The rank of $\psi(t^1\otimes\cdots\otimes t^k)$ is therefore at most $d_\varepsilon n^{2+\varepsilon}$ for some constant $d_\varepsilon \in \mathbb{N}$. We conclude that

the rank of $\psi(\mathrm{T}_{n,\gamma}(C_{\ell+2}))$ is at most $c_\varepsilon d_\varepsilon n^{\ell-1+2+2\varepsilon} = c_\varepsilon d_\varepsilon n^{\ell+1+2\varepsilon}$, and thus $\mathrm{R}(\mathrm{T}_{n,\gamma}(C_{\ell+2})) = \mathcal{O}(n^{\ell+1+\varepsilon})$ for any $\varepsilon > 0$.

Symmetrizing $\mathrm{T}_{n,\gamma}(C_k)$ cyclically gives us a balanced cycle tensor, as follows:

$$\mathrm{T}_{n^{k-1}\lfloor n^\gamma\rfloor}(C_k) \cong \bigotimes_\pi \pi \cdot \mathrm{T}_{n,\gamma}(C_k),$$

where π goes over all powers of the cyclic permutation $(12\cdots k)$, and π acts by permuting the tensor legs. Let $\varepsilon > 0$ and let $\gamma < \alpha$. Then, by submultiplicativity of tensor rank,

$$\mathrm{R}(\mathrm{T}_{n^{k-1}\lfloor n^\gamma\rfloor}(C_k)) \leq \mathrm{R}(\mathrm{T}_{n,\gamma}(C_k))^k \leq c_{\varepsilon,\gamma}^k n^{(k-1+\varepsilon)k}.$$

Then, by Lemma 2.2,

$$\omega(\mathrm{T}(C_k)) \leq \frac{\log_n(c_{\varepsilon,\gamma}^k n^{(k-1+\varepsilon)k})}{\log_n(n^{k-1}\lfloor n^\gamma\rfloor)} \leq \frac{\log_n c_{\varepsilon,\gamma}^k + (k-1+\varepsilon)k}{k-1+\gamma-o(1)}.$$

Letting $n \to \infty$, $\varepsilon \to 0$, $\gamma \to \alpha$ gives

$$\omega(\mathrm{T}(C_k)) \leq \frac{k}{k-1+\alpha}(k-1) = k - \alpha\Big(1 + \frac{1-\alpha}{k-1+\alpha}\Big),$$

finishing the proof. $\qquad\square$

Remark 3.10. We can naturally define $\omega(\gamma_1,\ldots,\gamma_k)$ by extending the definition in (5). An interesting intermediate result in the above proof of Theorem 3.9 is that for any $k \geq 3$ and any $0 < \gamma < \alpha$ we have $\omega(1,1,\ldots,1,\gamma) = k-1$. The standard flattening argument implies that this bound is optimal. Also, the observation in Remark 3.8 applied to $\mathrm{T}_{n,\gamma}(C_k)$ implies that the decompositions achieving the exponent $k-1+\varepsilon$ must have close to maximal local ranks, and thus the surgery method cannot be improved by taking into account local rank information.

Summarizing, the following table contains the best bounds on the exponent of odd cycles $\omega_k = \omega(\mathrm{T}(C_k))$ for some small odd k. From $k = 11$ onwards, Theorem 3.9 gives the best upper bound. This bound converges to $k - \alpha$ when we let $k \to \infty$.

k	ω_k		reference
	lower	upper	
3	2	2.3728639	[LG14]
5	4	4.6031719	[BCZ17]
7	6	6.6511249	[BCZ17]
9	8	8.6715848	Theorem 3.9
11	10	10.676522	Theorem 3.9
13	12	12.679854	Theorem 3.9

3.3 Covering and distilling

In some cases, we have another method for obtaining upper bounds on $\omega(\mathrm{T}(G))$ which gives weaker results than the tensor surgery upper bounds above but which is conceptually easier. The idea is to cover the graph G with triangles, which cost ω each and use distillation to remove unwanted edges.

For example, for $k = 5$, the distillation result (4) says that asymptotically one copy of $\mathrm{T}(C_5)$ can be restricted to the tensor product of two copies of $\mathrm{T}(5) = \sum_{i \in \{0,1\}} b_i^{\otimes 5}$. Covering the complete graph K_5 with 10 triangles, gives, with subscripts denoting tensor leg positions,

$$\mathrm{T}_{n^3}(K_5) \cong \bigotimes_{\substack{G \subseteq K_5: \\ G \cong C_3}} \mathrm{T}_n(C_3)_{V(G)} \otimes (1 \otimes \cdots \otimes 1)_{[k] \setminus V(G)},$$

where the tensor product is over subgraphs G of K_5 isomorphic to C_3. We can view $\mathrm{T}_{n^3}(K_5)$ as the tensor product of $\mathrm{T}_{n^3}(C_5)$ and a permuted copy of $\mathrm{T}_{n^3}(C_5)$. Distilling a unit tensor $\mathrm{T}_{\Omega(n^{2 \cdot 3 - \varepsilon})}(5)$ from one of these copies (Lemma 2.3) gives

$$\mathrm{T}_{n^3}(C_5)^{\oplus \Omega(n^{2 \cdot 3 - \varepsilon})} \cong \mathrm{T}_{n^3}(C_5) \otimes \mathrm{T}_{\Omega(n^{2 \cdot 3 - \varepsilon})}(5) \leq \mathrm{T}_{n^3}(K_5).$$

By the asymptotic sum inequality for cycles (Proposition 27 in [BCZ17]) we obtain the inequality $\omega(\mathrm{T}(C_5)) \leq (10\omega - 2 \cdot 3)/3$ which is at most 5.90955 by Le Gall's upper bound on ω.

A variation on the above idea is to cover the cycle C_k by unbalanced triangles with edge-multiplicities $(1, 1, \alpha)$, which cost 2 each, and then distil a k-cycle with multiplicity α. This yields $\omega(\mathrm{T}(C_k)) \leq k - \alpha$.

## 4 Tensor surgery on general graphs and hypergraphs

In this final section we want to illustrate tensor surgery on general graphs and hypergraphs. The first example shows that tensor surgery on a graph might involve absorbing a virtual hyperedge. The second example is an example of general hypergraph surgery. We believe that the bounds in this section cannot be obtained by using only the covering and distilling technique mentioned at the end of the previous section.

4.1 The dome tensor

In both examples we use the following hypergraph tensor, of which we will first establish some properties.

We define $\mathrm{dome}_{k,\ell}$ to be the following hypergraph on 4 vertices with

multi-edges

where k and ℓ denote edge-multiplicities.

Lemma 4.1. *We have* $3 \leq \omega(\mathrm{T}(\mathrm{dome}_{1,1})) \leq 3\omega/2$.

Proof. The lower bound $3 \leq \omega(\mathrm{T}(\mathrm{dome}_{1,1}))$ is obtained by grouping the black vertices together and taking the corresponding flattening of $\mathrm{T}(\mathrm{dome}_{1,1})$. For the upper bound, first observe that the exponent of the tensor $\mathrm{T}(G)$ corresponding to the graph G given by

is at most 3ω, since $\mathrm{T}(G)$ can be obtained by combining three copies of $\mathrm{T}(C_3)$. The distillation result (3) says that $\omega(\mathrm{T}(C_3), \mathrm{T}(3))^{-1} = 2$. This means that for any $\varepsilon > 0$ we can restrict $\mathrm{T}_n(C_3)$ to $\mathrm{T}_{\Omega(n^{2-\varepsilon})}(T(3))$ (Lemma 2.3). Applying this observation to the copy of C_3 that forms the base triangle in G gives that for any $\varepsilon > 0$ the tensor $\mathrm{T}_{\Omega(n^{2-\varepsilon})}(\mathrm{dome}_{1,1})$ has rank $\mathcal{O}(n^{3\omega+\varepsilon})$. Therefore, $\omega(\mathrm{T}(\mathrm{dome}_{1,1})) = 3\omega$. $\qquad\square$

Lemma 4.2. $\omega(\mathrm{T}(\mathrm{dome}_{1,4})) = 12$.

Proof. The lower bound $3 \cdot 4 \leq \omega(\mathrm{T}(\mathrm{dome}_{1,4}))$ is obtained by grouping the black vertices together and taking the corresponding flattening of $\mathrm{T}(\mathrm{dome}_{1,4})$. We prove the upper bound by proving that $\omega(\mathrm{T}(\mathrm{dome}_{1,4})^{\otimes 2}) \leq 24$. We will do this by following the strategy of the proof of Lemma 4.1. Recall that $\omega(\langle n, n^4, n^4 \rangle)$ equals $\omega(1, 4, 4) = 4\omega(\frac{1}{4}, 1, 1)$ (Lemma 2.4) and this number equals $4 \cdot 2 = 8$, since the dual exponent of matrix multiplication α is at least 0.3029805 which is strictly more than $\frac{1}{4}$ (see (6)). Therefore, the exponent of the tensor $\mathrm{T}(G)$ corresponding to the graph G given by

is at most $3\omega(1, 4, 4) = 24$. For any $\varepsilon > 0$ we can restrict $\mathrm{T}_n(C_3)$ to $\mathrm{T}_{\Omega(n^{2-\varepsilon})}(3)$ (Lemma 2.3). So for any $\varepsilon > 0$ the tensor $\mathrm{T}_{\Omega(n^{2-\varepsilon})}(\mathrm{dome}_{1,4})$ has rank $\mathcal{O}(n^{24+\varepsilon})$, which means that the inequality $\omega(\mathrm{T}(\mathrm{dome}_{1,4})^{\otimes 2}) \leq 24$ holds (Lemma 2.2). $\qquad\square$

4.2 Tensor surgery for graphs with hypergraph insertion

The aim of the first example is to show how tensor surgery on a graph may involve the absorption of a virtual hyperedge. Let G be the multigraph


where the numbers denote edge-multiplicity. Grouping the white vertices together and grouping the black vertices together shows that the size of a max-cut is at least 32. Therefore, $\omega(\mathrm{T}(G)) \geq 32$. On the other hand, one can cover the 5-cycle on the left at cost ω_5 and the remaining edges at cost 1 each, which implies that $\omega(\mathrm{T}(G)) \leq \omega(\mathrm{T}(C_5)) + 28$. Therefore, by Theorem 3.5 if $\omega = 2$, then $\omega(\mathrm{T}(G)) = 32$. We will now prove this bound independently of ω being 2.

Proposition 4.3. $\omega(\mathrm{T}(G)) = 32$.

Proof. It remains to show the upper bound. We start off with the rectangular matrix multiplication tensor $\langle n, n^4, n^4 \rangle$ at cost $\omega(1, 4, 4) = 4\,\omega(\frac{1}{4}, 1, 1) = 8$ (by Lemma 2.4 and since $\frac{1}{4} < \alpha$), and, viewing it as a triangle graph

split up one of the low-dimension vertices into three vertices such that the resulting tensor corresponds to the following graph:


We then insert a new vertex and edges with multiplicity 8 as follows:

Since the rank of a tensor in $\mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^{n^3}$ is at most n^2, the linear map which splits up the vertex and inserts the new vertex together with the appropriate edges with multiplicity 8 has cost at most the cost of creating the tensor corresponding to the hypergraph $\mathrm{T}(\mathrm{dome}_{1,4})^{\otimes 2}$ of Lemma 4.2. Thus, $\omega(\mathrm{T}(G)) \leq 4\,\omega(\frac{1}{4}, 1, 1) + 2\,\omega(\mathrm{T}(\mathrm{dome}_{1,4})) \leq 4 \cdot 2 + 2 \cdot 12 = 32$. $\qquad\square$

4.3 Tensor surgery for hypergraphs

In the second example we will be inserting a hypergraph into a hypergraph. Define H as the hypergraph

Proposition 4.4. *We have $6 \leq \omega(\mathrm{T}(H)) \leq 6\omega/2$.*

Proof. The lower bound follows from grouping the white vertices together and grouping the black vertices together, and taking the corresponding flattening. For the upper bound, we start off with the dome $\mathrm{dome}_{1,1}$

We split one of the vertices in the hyperedge, as follows

and insert the remaining vertices and edges as to obtain the goal tensor.

We see that the combined cost of splitting the vertex and inserting the vertices and edges is at most $\omega(\mathrm{T}(\mathrm{dome}_{1,1}))$ which is at most $3\omega/2$ (Lemma 4.1). We conclude that the inquality $\omega(\mathrm{T}(G)) \leq 2\,\omega(\mathrm{T}(\mathrm{dome}_{1,1})) \leq 6\omega/2$ holds. $\qquad\square$

Of course, by replacing $\mathrm{dome}_{1,1}$ by $\mathrm{dome}_{1,4}$ one can obtain an exact result like in Proposition 4.3.

One of the reviewers observed the following. The upper bound of Proposition 4.4 can also be obtained by covering by cycle graphs and distillation in a way similar to Lem. 4.1: Number the vertices of H left to right, top to bottom. Combine the 5-cycles 12463, 12765 and 3-cycles 135, 247, then distill 3-cycles 356 and 467 to obtain hyperedges. The resulting hypergraph is the hypergraph H doubled.

References

[AFLG15] Andris Ambainis, Yuval Filmus, and François Le Gall. Fast matrix multiplication: limitations of the Coppersmith-Winograd method. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 585–593. ACM, New York, 2015.

[BCS97] Peter Bürgisser, Michael Clausen, and M. Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren Math. Wiss.* Springer-Verlag, Berlin, 1997.

[BCZ17] Harry Buhrman, Matthias Christandl, and Jeroen Zuiddam. Nondeterministic Quantum Communication Complexity: the Cyclic Equality Game and Iterated Matrix Multiplication. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, pages 24:1–24:18, 2017.

[BI11] Peter Bürgisser and Christian Ikenmeyer. Geometric complexity theory and tensor rank. In *STOC'11—Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 509–518. ACM, New York, 2011.

[Blä03] Markus Bläser. On the complexity of the multiplication of matrices of small formats. *J. Complexity*, 19(1):43–60, 2003.

[CCD+10] Lin Chen, Eric Chitambar, Runyao Duan, Zhengfeng Ji, and Andreas Winter. Tensor rank and stochastic entanglement catalysis for multipartite pure states. *Phys. Rev. Lett.*, 105(20):200501, 2010.

[CM96] Pierre Comon and Bernard Mourrain. Decomposition of quantics in sums of powers of linear forms. *Signal Processing*, 53(2):93–107, 1996.

[CU12] Henry Cohn and Christopher Umans. Fast matrix multiplication using coherent configurations. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1074–1087. SIAM, Philadelphia, PA, 2012.

[CVZ16]   Matthias Christandl, Péter Vrana, and Jeroen Zuiddam. Asymptotic tensor rank of graph tensors: beyond matrix multiplication. *arXiv*, 2016. `arXiv:1609.07476`.

[CW90]   Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. *J. Symbolic Comput.*, 9(3):251–280, 1990.

[dG78]   Hans F. de Groote. On varieties of optimal algorithms for the computation of bilinear mappings. I. The isotropy group of a bilinear mapping. *Theoret. Comput. Sci.*, 7(1):1–24, 1978.

[DKW11]   Jan Draisma, Eyal Kushilevitz, and Enav Weinreb. Partition arguments in multiparty communication complexity. *Theoret. Comput. Sci.*, 412(24):2611–2622, 2011.

[Hås90]   Johan Håstad. Tensor rank is NP-complete. *J. Algorithms*, 11(4):644–654, 1990.

[HK71]   John E. Hopcroft and Leslie R. Kerr. On minimizing the number of multiplications necessary for matrix multiplication. *SIAM J. Appl. Math.*, 20:30–36, 1971.

[Lan06]   Joseph M. Landsberg. The border rank of the multiplication of 2×2 matrices is seven. *J. Amer. Math. Soc.*, 19(2):447–459, 2006.

[Lan12]   Joseph M. Landsberg. *Tensors: geometry and applications*, volume 128 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2012.

[LG12]   François Le Gall. Faster algorithms for rectangular matrix multiplication. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science—FOCS 2012*, pages 514–523. IEEE Computer Soc., Los Alamitos, CA, 2012.

[LG14]   François Le Gall. Powers of tensors and fast matrix multiplication. In *ISSAC 2014—Proceedings of the 39th International Symposium on Symbolic and Algebraic Computation*, pages 296–303. ACM, New York, 2014.

[LO13]   Joseph M. Landsberg and Giorgio Ottaviani. Equations for secant varieties of Veronese and other varieties. *Ann. Mat. Pura Appl. (4)*, 192(4):569–606, 2013.

[LO15]   Joseph M. Landsberg and Giorgio Ottaviani. New lower bounds for the border rank of matrix multiplication. *Theory Comput.*, 11:285–298, 2015.

[PS05]   Lior Pachter and Bernd Sturmfels, editors. *Algebraic Statistics for Computational Biology*. Cambridge University Press, Aug 2005.

[Shi16]   Yaroslav Shitov. How hard is the tensor rank? *arXiv preprint arXiv:1611.01559*, 2016.

[Smi13]   A. V. Smirnov. The bilinear complexity and practical algorithms for matrix multiplication. *Zh. Vychisl. Mat. Mat. Fiz.*, 53(12):1970–1984, 2013.

[SS16]   Marcus Schaefer and Daniel Stefankovic. The Complexity of Tensor Rank. *arXiv preprint arXiv:1612.04338*, 2016.

[Str69]   Volker Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13(4):354–356, 1969.

[Str83]   Volker Strassen. Rank and optimal computation of generic tensors. *Linear Algebra Appl.*, 52/53:645–685, 1983.

[Str87]   Volker Strassen. Relative bilinear complexity and matrix multiplication. *J. Reine Angew. Math.*, 375/376:406–443, 1987.

[Str88]   Volker Strassen. The asymptotic spectrum of tensors. *J. Reine Angew. Math.*, 384:102–152, 1988.

[Str91]   Volker Strassen. Degeneration and complexity of bilinear maps: some asymptotic spectra. *J. Reine Angew. Math.*, 413:127–180, 1991.

[Str05]   Volker Strassen. Komplexität und Geometrie bilinearer Abbildungen. *Jahresber. Deutsch. Math.-Verein.*, 107(1):3–31, 2005.

[VC15]   Péter Vrana and Matthias Christandl. Asymptotic entanglement transformation between W and GHZ states. *J. Math. Phys.*, 56(2):022204, 12, 2015. arXiv:1310.3244.

[VC17]   Péter Vrana and Matthias Christandl. Entanglement Distillation from Greenberger–Horne–Zeilinger Shares. *Commun. Math. Phys.*, 352(2):621–627, 2017. arXiv:1603.03964.

[Win71]   Shmuel Winograd. On multiplication of 2×2 matrices. *Linear Algebra Appl.*, 4(4):381–388, 1971.