Introduction

The following is an exposition of a recent paper [IMW17] which proves apart from the NP-hardness of KRONECKER, the existence of superpolynomially many (pseudo) ”exceptional” Kronecker coefficients. These coefficients are relevant for the GCT program which is the motivation for the renewed interest in this topic. Though the applications of the paper are clearly to complexity theory and representation theory, the tools are purely combinatorial and this exposition thus assumes no background.

Preliminaries

Definition 1. *Partition - A partition λ is defined to be a finite non-increasing sequence of positive integers $\lambda = (\lambda_1, \cdots, \lambda_l)$. We define height $ht(\lambda) = l$. Define the conjugate $\lambda^T := (\lambda_i^T) =$ number of boxes in i^{th} column.*

Example -
If $\lambda = (3, 2)$ the diagram is 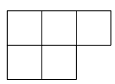 and the conjugate i.e. $\lambda^T = (3, 2, 1, 1)$ is 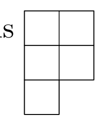

- We'll do something very interesting. We'll prove results about the Kronecker coefficient without even requiring to define it! Obviously we'll need to borrow a (simple) lemma but the rest would be self contained.
- The problem in working directly with the Kronecker coefficient, is that it doesn't have a combinatorial definition. In fact, finding one is a long standing open problem in algebraic combinatorics posed by Murnaghan in 1930s.
- So how do we work with it if we don't have a definition? A clever idea - sandwich it between 2 other quantities that have both a combinatorial and a representation theoretic interpretation.
- The interested reader may read this for the definition of kronecker coefficient and the connections to GCT.

Given a (finite) point set $P \subset \{0, ..., r-1\}^3$, let $x_P(i)$, be the number of points in P with the x-coordinate i. We call $x_P = (x_P(0), \cdots, x_P(r-1))$ the x-marginal of P. Similarly define the y-marginal y_P and the z-marginal z_P . The triple (x_P, y_P, z_P) is called the marginals of P.

Definition 2. *Pyramid - A point set P as above is called a pyramid if it's downward closed as in - $\forall (x, y, z) \in P$ and any point $0 \le x' \le x$, $0 \le y' \le y$, $0 \le z' \le z$ $(x', y', z') \in P$*

Definition 3. $t^{\lambda}_{\mu,\pi}$ *- Number of P such that $(x_P, y_P, z_P) = (\lambda^T, \mu^T, \pi^T)$*

Observation - For a marginal to be denoted by a partition, $x_P(0) \ge x_P(1) \ge \cdots \ge x_p(r-1)$ and similarly for y_P and z_P. But if P is a pyramid, then this condition is automatically satisfied.

Definition 4. $p^{\lambda}_{\mu,\pi}$ *- Number of pyramid-P such that* $(x_P, y_P, z_P) = (\lambda^T, \mu^T, \pi^T)$

Clearly then from the definitions, $p^{\lambda}_{\mu,\pi} \leq t^{\lambda}_{\mu,\pi}$

Now we state without proving the lemma which will let us forget about the Kronecker coefficient and work with t and p. This is proved by interpreting t, p in a representation theory fashion.

Lemma 1. *[IMW17]* $p^{\lambda}_{\mu,\pi} \leq k^{\lambda}_{\mu,\pi} \leq t^{\lambda}_{\mu,\pi}$

NP-hardness

The aim is now to prove that the decision problem KRONECKER, which takes (λ, μ, π) in unary as input and decides whether $k^{\lambda}_{\mu,\pi} > 0$, is NP-hard. The strategy is now very simple.

1. Find a sufficient condition on the triple (λ, μ, π) such that every point-set is a pyramid.
2. Show that such triples exist. Thus, $k^{\lambda}_{\mu,\pi} = t^{\lambda}_{\mu,\pi}$.
3. Apply the result of [BLG01] which says that checking the positivity of $t^{\lambda}_{\mu,\pi}$ is NP-hard even when (λ, μ, π) is of the above-type.

The sufficient condition

The key idea is to look at point sets that minimize a function f such that f is a function only of the marginals and attains the minima only at pyramids.

The function we choose is the barycenter i.e.

$$b(P) = b(x_P, y_P, z_P) = \sum_{(x,y,z)\in P} x + y + z = \sum_i i\,(x_P(i) + y_P(i) + z_P(i))$$

Define a r-simplex as $P_r = \{(x, y, z)|x + y + z \leq r - 1\}$. For any n, define $r(n) = max\{r \mid |P_r| = \binom{r+2}{3} \leq n\}$. It's clear from definition of the r-simplex that for any point set Q of size n, $b(Q) \geq b(P_{r(n)}) + r(n)(n - |P_r(n)|) := p(n)$ Equality holds if only if $P_{r(n)} \subset Q \subset P_{r(n)+1}$

Definition 5 (Simplex-like)**.** *A triple of partitions* (λ, μ, π) *is called simplex-like if each of them has atmost* $r + 1$ *columns and*

$$b(\lambda, \mu, \pi) = \sum_{i=0}^{r} i\lambda_i^T + \sum_{i=0}^{r} i\mu_i^T + \sum_{i=0}^{r} i\pi_i^T = p(n)$$

Theorem 2. *If* (λ, μ, π) *is simplex-like, then* $p^{\lambda}_{\mu,\pi} = t^{\lambda}_{\mu,\pi}$.

Proof. Let Q be of size n and have marginals (λ, μ, π). $b(Q) = b(\lambda, \mu, \pi) = p(n)$. Thus, from the above observation, $P_{r(n)} \subset Q \subset P_{r(n)+1}$ and in particular is a pyramid. □

The [BLG01] construction

The work of [BLG01] looks at at a subset of these simplex-like partition triples for which they prove the NP-hardness result. Their specific construction is to take P_{2r} and add $r+1$ points such that $\lambda^T = \phi^T + (d_0, \cdots, d_{2r})$ and $\mu^T = \pi^T = \phi^T + (1^{r+1}, 0^r)$ such that $\sum_k d_k = r+1$ and $\sum_k k(d_k) = r(r+1)$.

Theorem 3. *[BLG01] KRONECKER is NP-hard even when (λ, μ, π) is of this restricted type as above.*

Complexity of $t^{\lambda}_{\mu,\pi}$ and relations with $k^{\lambda}_{\mu,\pi}$

We know certain that the Kronecker coefficient can be computed efficiently for certain "natural" sub-classes of partition triples. This begs us to consider if positivity of t for the same classes can also be computed efficiently.

Family	$k^{\lambda}_{\mu,\pi}$ [Known]	$t^{\lambda}_{\mu,\pi}$ [IMW17]
Littlewood Richardson coefficient*	P [KT99, MNS12]	Always > 0
Constant height	P	P
λ is a hook	$\#P$ [Bla17]	P
Rectangular	Conj - [Mul07]	P

* - A recent work has generalized this case to Littlewood-Richardson polynomials - [AA17].

The paper defines another quantity $\widetilde{t}^{\lambda}_{\mu,\pi}$ as the number of hypergraphs of type - (λ, μ, π) with certain properties. It then shows that $\widetilde{t}^{\lambda}_{\mu,\pi} > 0 \iff t^{\lambda}_{\mu,\pi} > 0$. I find that approach unnecessary, at least for the cases 1 and 3, and will thus prove it directly using t.

λ is a hook

Let $\lambda = (D-k+1, 1^{k-1})$ i.e. $\lambda^T = (k, 1^{D-k})$. and μ, π be any partitions of D. Let $ht(\mu^T) = h_1, ht(\pi^T) = h_2$. We need to decide if there exists a point set P of type (λ, μ, π). Say, you're given a point set P with y, z labels such that $\vec{y_P} = \mu$ and $\vec{z_P} = \pi$.

To formalize this, we can identify $P \cong [D]$ by labelling the vertices but we define it generally to avoid confusion. Define $l : [D] \to [r] \times [r]$. l is of type (μ, π) if $\forall i \ \sum_j |l^{-1}(i,j)| = \mu_i^T$ and $\forall j \ \sum_j |l^{-1}(i,j)| = \pi_j^T$ We write such a label as $l_{\mu,\pi}$. What necessary and sufficient conditions must $l_{\mu,\pi}$ posses to be able to extend to a *valid* pointset?

Define an equivalence relation $v_1 \sim v_2$ iff $l(v_1) = l(v_2)$. Let $|l|$ denote the number of equivalence classes which is the same as the size of its image. Another way of representing l is by giving the list equivalence classes (i,j) and then giving the number of vertices in each class. Thus l can be represented as a $h_1 \times h_2$ matrix (n_{ij}) where n_{ij} is the number of vertices with label (i,j). Moreover, given any such matrix (of non-negative integers) such that $\sum_{i,j} n_{ij} = D, \sum_j n_{ij} = \mu_i^T, \sum_i n_{ij} = \pi_j^T$ we have a labelling $l_{\mu,\pi}$ where $|l|$ is the number of non-zero entries .

Lemma 4. $t^{\lambda}_{\mu,\pi} > 0 \iff \exists \ l_{\mu,\pi}$ *such that* $|l| \geq k$

Proof. $\Rightarrow$ Since $\lambda^T = (k, 1^{D-k})$ we need k points with x-coordinate 0. But if $|l| < k$ then by pigeonhole principle, any set of k points consists of 2 such that $l(v_1) = l(v_2)$ and thus they have the same x,y,z coordinate.
$\Leftarrow$ Given such a l, take any k points with distinct l-values. Assign their x-coordinate to be 0. For the rest, assign them arbitrarily with x-coordinates $2, 3 \cdots D - k + 1$. $\square$

Define $G_{\mu,\pi}$ which has $V = \{s, t, a_1, \cdots a_{h_1}, b_1, \cdots, b_{h_2}\}$, $E = E_s \cup E_m \cup E_t$ where,
$E_s = \{(s, a_i)|\ i \in [h_1],\ \}\ \ c(s, a_i) = \mu_i^T$
$E_m = \{(a_i, b_j)| i \in [h_1], j \in [h_2]\}\ \ c(a_i, b_j) = 1$
$E_t = \{(b_j, t)|\ j \in [h_2]\}\ \ c(b_j, t) = \pi_i^T$

A picture is worth a thousand word, so here goes.

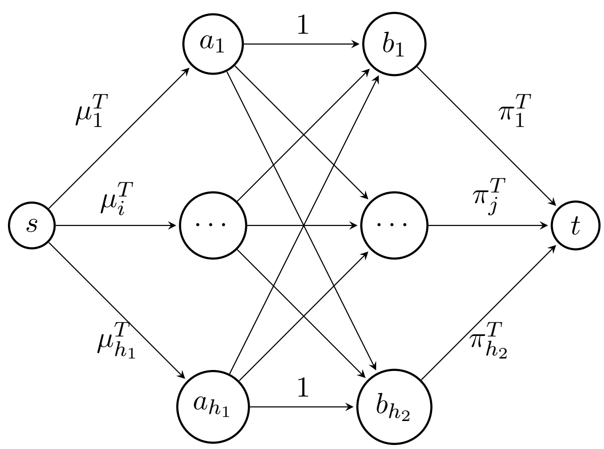

Theorem 5. $Maxflow(G_{\mu,\pi}) \geq k \iff \exists\ l_{\mu,\pi}$ *such that* $|l| \geq k$

Proof. Given a l, it's easy to construct a flow. Just assign $f(a_i, b_j) = 1$ if $(i, j) \in Im(l)$. Assign the flows to the source/sink edges appropriately to satisfy the constraints. This gives a valid flow as $f(s, a_i) = \sum_j \mathbb{I}\left(|l^{-1}(i, j)|\right) \leq \sum_j |l^{-1}(i, j)| = \mu_i^T$ where $\mathbb{I}(x) = 1$ if $x > 0$, else, 0. Similar relation holds for $f(b_j, t)$.

Given a maxflow f (of size $\geq k$), we get us a subset of our equvalence classes in l. Now to get their sizes, draw a residual graph with no constraints on the internal edges i.e ,

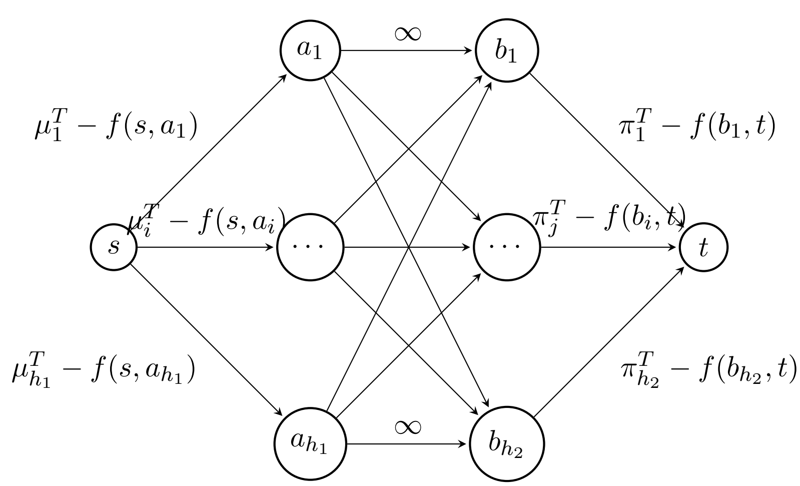

Solve the maxflow for this and say you get f'. Now, it is easy to see(directly or by Maxflow-mincut theorem) that the total flow, i.e $\sum_i (f(s,a_i) + f'(s,a_i)) = D$. Define $n_{ij} = f(a_i,b_j) + f'(a_i,b_j)$. Clearly, $\sum_{i,j} n_{ij} = D$ and since $|f| > k$ there are atleast k (i,j) such that $n_{ij} > 0$. and this matrix (n_{ij}) thus defines a l with at least k equivalence classes. □

A bit of Rep theory and the GCT connection

Defining $k^{\lambda}_{\mu,\pi}$

- This section is meant to highlight the motivation for studying this problem and is not critical for the rest of the exposition. So feel free to skip it.
- Given a group G and a vector space V, a representation of G on V is a group homomorphism $\rho : G \to GL(V)$ i.e. if $V \cong k^n$ for soem field k, then $\rho(g) \in GL_n(k)$ such that $\rho(gh) = \rho(g)\rho(h)$.
- $W \subset V$ is called subrepresentation if $\rho(g)W \subset W \quad \forall g \in G$. V is irreducible if its has no non-trivial subrepresentations.
- A fundamental theorem of Maschke(for finite groups) and Weyl (for more general *reductive* groups) is that every representation W of a finite (reductive, in general) group G can be decomposed into irreducible representations i.e. $W = \bigoplus_i a_i V_i$
- If (V,ρ) and (W,δ) are representations of G then so is $(V \otimes W, \rho \otimes \delta)$.
- For the symmetric group S_n its irreducible representations V_λ are indexed by partitions of n, i.e. λ such that $|\lambda| = n$.
- Decomposing the tensor product of 2 of these, we get the definition of the kronecker coefficient. $V_\mu \otimes V_\pi = \bigoplus_\lambda k^{\lambda}_{\mu,\pi} V_\lambda$

GCT approach - A (extremely) high level view

- Valiant defined an algebraic computation model and therefore, an analog of boolean complexity classes P, NP called VP and VNP.
- $\mathsf{VP}_{\mathsf{ws}} \subset \mathsf{VP}$ where *ws* stands for weakly skew. (You may omit this slight technicality for now) . Valiant conjectured that $\mathsf{VP}_{\mathsf{ws}} \neq \mathsf{VNP}$
- $\mathsf{P} \neq \mathsf{NP} \implies \mathsf{VP} \neq \mathsf{VNP} \implies \mathsf{VP}_{\mathsf{ws}} \neq \mathsf{VNP}$ but the other direction is not yet known.
- These algebraic classes are defined in terms of families of polynomials. The determinant in $\mathsf{VP}_{\mathsf{ws}}-$complete and the permanent is $\mathsf{VNP}-$complete. So, roughly we are saying that the permanent of a $n \times n$ matrix cannot be written as the determinant of a $p(n) \times p(n)$ sized matrix for any polynomial p.

- GCT aims to attack this problem using tools from algebraic geometry, representation theory and geometric invariant theory. It associates with these complexity classes certain representations of GL_n say $\mathsf{VP}_{\mathsf{ws}} \sim V_d$ and $\mathsf{VNP} \sim V_p$
- GCT reformulation - [MS01] - $\mathsf{VNP} \subseteq \overline{\mathsf{VP}_{\mathsf{ws}}} \implies V_p \subseteq V_d$. The bar signifies a (Zariski) closure but let's not get into that.
- Breaking into irreps, say $V_p = \bigoplus_\lambda a_\lambda V_\lambda$ and $V_d = \bigoplus_\lambda a'_\lambda V_\lambda$. An easy consequence of the Schur's lemma is that $V_p \subseteq V_d \implies a_\lambda \leq a'_\lambda \ \forall \ \lambda$
- Thus, finding a λ that violates the inequality would prove the conjecture.

References

[AA17] Alexander Yong Anshul Adve, Colleen Robichaux. Vanishing of littlewood-richardson polynomials is in p. *CoRR*, abs/1708.04228, 2017.

[Bla17] Jonah Blasiak. Kronecker coefficients for one hook shape. *Seminaire Lotharingien de Combinatoire*, 2017.

[BLG01] S. Brunetti, A. Del Lungo, and Y. Gerard. On the computational complexity of reconstructing three-dimensional lattice sets from their two-dimensional x-rays. *Linear Algebra and its Applications*, 339(1):59 – 73, 2001.

[IMW17] Christian Ikenmeyer, Ketan D. Mulmuley, and Michael Walter. On vanishing of kronecker coefficients. *Computational Complexity*, 26(4):949–992, 2017.

[KT99] Allen Knutson and Terence Tao. The honeycomb model of $\mathrm{GL}_n(\mathbf{C})$ tensor products. I : Proof of the saturation conjecture. *J. Amer. Math. Soc.*, 12:1055–1090, 1999.

[MNS12] Ketan D. Mulmuley, Hariharan Narayanan, and Milind Sohoni. Geometric complexity theory iii: on deciding nonvanishing of a littlewood–richardson coefficient. *Journal of Algebraic Combinatorics*, 36(1):103–110, Aug 2012.

[MS01] Ketan Mulmuley and Milind A. Sohoni. Geometric complexity theory I: an approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.

[Mul07] Ketan Mulmuley. Geometric complexity theory VI: the flip via saturated and positive integer programming in representation theory and algebraic geometry. *CoRR*, abs/0704.0229, 2007.