

THE HONEYCOMB MODEL OF $GL_n(\mathbb{C})$ TENSOR PRODUCTS I: PROOF OF THE SATURATION CONJECTURE

ALLEN KNUTSON AND TERENCE TAO

1. THE SATURATION CONJECTURE

A very old and fundamental question about the representation theory of $GL_n(\mathbb{C})$ is the following:

For which triples of dominant weights λ, μ, ν does the tensor product $V_\lambda \otimes V_\mu \otimes V_\nu$ of the irreducible representations with those high weights contain a $GL_n(\mathbb{C})$ -invariant vector?

Another standard, if less symmetric, formulation of the problem above replaces V_ν with its dual, and asks for which ν is V_ν^* a constituent of $V_\lambda \otimes V_\mu$. In this formulation one can without essential loss of generality restrict to the case that λ, μ , and ν^* are polynomial representations, and rephrase the question in the language of Littlewood-Richardson coefficients; it asks for which triple of partitions λ, μ, ν^* is the Littlewood-Richardson coefficient $c_{\lambda\mu}^{\nu^*}$ positive.

It is not hard to prove (as we will see later in this introduction) that the set of such triples (λ, μ, ν) is closed under addition, so forms a monoid. In this paper we prove that this monoid is *saturated*, i.e. that for each triple of dominant weights (λ, μ, ν) ,

$$(V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu})^{GL_n(\mathbb{C})} > 0 \text{ for some } N > 0 \implies (V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})} > 0.$$

This is of particular interest because Klyachko has recently given an answer¹ to the general question above, which in one direction was only asymptotic [Kl]:

If $V_\lambda \otimes V_\mu \otimes V_\nu$ has a $GL_n(\mathbb{C})$ -invariant vector, then λ, μ, ν satisfy a certain system of linear inequalities derived from Schubert calculus (plus the evident linear equality that $\lambda + \mu + \nu$ be in the root lattice; in the L-R context this asks that the number of boxes

Received by the editors July 31, 1998 and, in revised form, February 25, 1999.

1991 *Mathematics Subject Classification*. Primary 05E15, 22E46; Secondary 15A42.

Key words and phrases. Honeycombs, Littlewood-Richardson coefficients, Berenstein-Zelevinsky patterns, Horn's conjecture, saturation, Klyachko inequalities.

The first author was supported by an NSF Postdoctoral Fellowship.

The second author was partially supported by NSF grant DMS-9706764.

¹Klyachko gives a finite set of inequalities, that *as a set* are necessary and sufficient for this asymptotic result. However, Chris Woodward has informed us that contrary to Klyachko's unproven claim in [Kl], the inequalities are not independent – not all of them determine facets of the cone. This will be the subject of inquiry of our second paper [Hon2].

in the partition ν^* is the number of boxes in λ and μ together). Conversely, if λ, μ, ν satisfy these inequalities, then there exists an integer N such that the tensor product $V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu}$ has a $GL_n(\mathbb{C})$ -invariant vector.

Our saturation result completes this converse, saying that Klyachko's inequalities completely characterize the monoid. The survey papers [F, Ze] point out another important consequence of these two results taken together: *Horn's conjecture* [H] from 1962, which gives a recursive system of inequalities, since the relevant Schubert calculus questions can be cast as lower-dimensional Littlewood-Richardson questions.

The main tool in this paper is the Berenstein-Zelevinsky cone [BZ, Ze], and in particular the BZ polytope associated to the triple (λ, μ, ν) , in which the number of lattice points is the corresponding Littlewood-Richardson coefficient. We use a new description of the BZ cone: the *honeycomb* model. (The reader who is willing to grant Appendix 1 does not need to absorb separately the definition of the BZ cone.) This is a special case of a general way of producing polyhedra that we dub *tinkertoy models*. This viewpoint gives us natural ways to interpret faces of the BZ polytope as associated to simpler tinkertoys. In addition, the Gel'fand-Cetlin system fits in this theory as associated to a 1-dimensional tinkertoy.

The essence of the proof is as follows. We introduce a way of indexing (real) points in the Berenstein-Zelevinsky cone by planar pictures called *honeycomb diagrams*; this identification is in Appendix 1. This rational polyhedral cone linearly projects to the space of triples (λ, μ, ν) of (real) dominant weights of $GL_n(\mathbb{C})$. By [BZ], the number of integral points (honeycomb diagrams whose vertices lie at points in the triangular lattice) in a fiber of this projection is the dimension $(V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})}$. In order to work conveniently with honeycomb diagrams, we introduce the seemingly richer notion of a honeycomb, since honeycombs can be seen to naturally fit into a polyhedral cone. Then we prove the somewhat technical Theorem 1 that honeycombs are characterized by their diagrams (whose linear structure is less apparent).

If for some large N we have

$$(V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu})^{GL_n(\mathbb{C})} > 0,$$

then the fiber over $(N\lambda, N\mu, N\nu)$ of this linear projection contains a lattice point and is thus nonempty. By rescaling we find that the fiber over (λ, μ, ν) is also nonempty. So the question comes down to showing that a nonempty fiber over an integral triple necessarily contains a lattice honeycomb. Equivalently, we want a way of deforming a nonlattice honeycomb with integral "boundary conditions" λ, μ, ν to a lattice honeycomb.

We do this by maximizing a linear functional, the "weighted perimeter",² on the polytope of honeycombs with given λ, μ, ν . This picks out an extremal honeycomb,³ the "largest lift", which we prove in Theorem 2 to have very nice properties if the

²One point easily missed is that an arbitrary choice is made in choosing this functional, making the subsequent construction noncanonical – but since we only seek an existence proof, this is not a problem.

³It was not a priori obvious that the lattice point we seek occurs as a vertex of the honeycomb polytope. In particular, not all the vertices are at lattice points (see Figure 18). Nor was it plain that a single functional could be used to pick them out uniformly for all λ, μ, ν . These facts are side consequences of the proof.

three weights are suitably generic. This theorem seems to be the useful one for studying honeycombs, and will play an equally important role in the next paper in this series [Hon2].

It is then straightforward to prove from its nice properties that the largest lift is integral. A continuity argument handles the case of nongeneric triples of weights. This ends the proof.

Not all of the framework presented in this paper is strictly necessary if one only wishes to prove the saturation conjecture. In the very nice paper [Bu] a streamlined version of our proof is presented, avoiding honeycombs in favor of the hive model⁴ (presented in Appendix 2), and in particular not requiring Theorem 1. However, one consequence of Theorem 1 is that honeycombs have a very important operation called *overlaying* which will be central for developments in later papers. In the next in this series we will use the overlaying operation to study which of Klyachko's inequalities are in fact essential [Hon2].

We thank Chris Woodward for pointing out that the saturation conjecture gives a new proof of the weak PRV conjecture for $GL_n(\mathbb{C})$, which states that $V_{w\lambda+v\nu}$ (for $w\lambda + v\nu$ in the positive Weyl chamber) is a constituent of $V_\lambda \otimes V_\mu$. (This “conjecture” is nowadays known to be true for all Lie groups [KMP].) In fact one can do better, and without using saturation; in section 4 there is a canonical honeycomb witnessing each instance of the long-proven “conjecture”, constructed by overlaying GL_1 -honeycombs.

We mention very briefly some connections to algebraic and symplectic geometry (much more can be found in [F, Ze]). By Borel-Weil, the space $(V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})}$ is the space of invariant sections of the (λ, μ, ν) line bundle on the product of three flag manifolds. Given two nonzero invariant sections, one of the (λ, μ, ν) line bundles and one of the (λ', μ', ν') , we can tensor them together to get an invariant section of the $(\lambda + \lambda', \mu + \mu', \nu + \nu')$ line bundle. The geometrical fact that the flag manifold (hence the product) is reduced and irreducible guarantees that this tensor product section is again nonzero; this is why the set of triples (λ, μ, ν) with invariant sections forms a monoid.

This same data is involved in defining a geometric invariant theory quotient of the product of three flag manifolds by the diagonal action of $GL_n(\mathbb{C})$. The space $(V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu})^{GL_n(\mathbb{C})}$ is the N th graded piece of the coordinate ring of this quotient space. (Klyachko's paper is a study of the semistability conditions that arise in performing this quotient.) The BZ counting result then says that this moduli space of triples of flags has the same Hilbert function as a certain toric variety, and our saturation result says that the (by definition ample) line bundle on this moduli space actually has sections. W. Fulton has shown us examples in which this line bundle is not very ample. If one had an explicit degeneration of the moduli space to the toric variety, one might be able to relate this non-very-amplitude to the existence of nonintegral vertices on the corresponding polytope of honeycombs.

The symplectic geometry connection then comes from the “GIT quotients are symplectic quotients” theorem [MFK, Chapter 8] (whose proof is essentially repeated in Klyachko's paper, in this special case). In this case, the corresponding symplectic quotient is the space of triples of Hermitian matrices with spectra $\lambda, \mu,$

⁴The first version of this paper required the reader to absorb both models and switched viewpoint back and forth. The paper [Bu] was inspired by that version, and took the approach of eliminating honeycombs, rather than hives as is done here.

and ν which sum to zero, modulo the diagonal action of $U(n)$:

$$\begin{aligned} & (Fl(\mathbb{C}^n)_\lambda \times Fl(\mathbb{C}^n)_\mu \times Fl(\mathbb{C}^n)_\nu) // GL_n(\mathbb{C}) \\ & \cong \{(H_\lambda, H_\mu, H_\nu) : \text{eigen}(H_\alpha) = \alpha, H_\lambda + H_\mu + H_\nu = 0\} / U(n). \end{aligned}$$

In particular, this identification shows directly that the existence of a $GL_n(\mathbb{C})$ -invariant vector in $V_\lambda \otimes V_\mu \otimes V_\nu$ implies the existence of a triple of Hermitian matrices with zero sum. The reverse implication exactly amounts to the saturation conjecture.

B. Sturmfels has pointed out that the “largest lift” construction can be interpreted as selecting a vertex of the *fiber polytope* [BS] of the projection from the cone of honeycombs to the cone of triples of dominant weights. Combining this idea with the hypothetical degeneration of the moduli space to the BZ toric variety, this suggests that we might be able to alternately interpret the largest lift as picking out a point in the *Chow* quotient [KSZ] of the product of three flag manifolds by the diagonal action of $GL_n(\mathbb{C})$.

We thank Anders Buch, Bill Fulton, Bernd Sturmfels, Greg Warrington, and Andrei Zelevinsky for careful readings and many cogent suggestions; we especially thank Bill for correcting a number of historical inaccuracies in the early versions.

We encourage the reader to get a feeling for honeycombs by playing with the honeycomb Java applet at

<http://www.alumni.caltech.edu/~allenk/java/honeycombs.html>.

2. TINKERTOYS AND THE HONEYCOMB MODEL

We first fix a few standard notations.

A weight of $GL_n(\mathbb{C})$ is a list of n integers, and is dominant if the list is weakly decreasing. So $GL_n(\mathbb{C})$'s root lattice is the hyperplane of lists whose sum is zero.

The “graphs” in this paper are rather nonstandard; for us, a directed graph Γ is a quadruple $(V_\Gamma, E_\Gamma, \text{head}, \text{tail})$ where the *head* and *tail* maps from the edges E_Γ to the vertices V_Γ may be only *partially defined* – the edges may be semi- or even fully infinite. In particular, any subset of the vertices and edges gives a subgraph, where the domains of definition of the *head* and *tail* maps are restricted to those edges that have their heads or tails in the subgraph.

For B a real vector space, let $\text{Rays}(B) := (B - \{0\}) / \mathbb{R}_+$ denote the space of rays coming from the origin. Each ray d is in a unique line $\mathbb{R} \cdot d$. Topologically $\text{Rays}(B)$ is a sphere, $S^{\dim B - 1}$.

2.1. Tinkertoys. We define a **tinkertoy** τ as a triple (B, Γ, d) consisting of a vector space B , a directed graph Γ (possibly with some zero- or one-ended edges), and a map $d : E_\Gamma \rightarrow \text{Rays}(B)$ assigning to each edge e a “direction” $d(e)$ in the sphere.⁵

Example 1. *Polytope tinkertoys.* Any polytope P in B gives a natural tinkertoy, just from the vertices, the edges (oriented arbitrarily), and their directions $d(e) := (\text{head}(e) - \text{tail}(e)) / \mathbb{R}^+ \in \text{Rays}(B)$. For example, each rectangle in \mathbb{R}^2 with

⁵A related, though much more restrictive, definition has recently appeared in [GZ], in a context quite related to the polytope tinkertoys in Example 1 following. In both cases, it is sort of unnatural to fix an orientation on the graph – really it is the “orientation times the direction” that comes into play.

edges aligned with the coordinate axes gives us the same polytope tinkertoy (up to isomorphism).

We define a **configuration h of a tinkertoy τ** as a function $h : V_\Gamma \rightarrow B$ assigning a point of B to each vertex, such that for each two-ended edge e

$$h(\text{head}(e)) - h(\text{tail}(e)) \in d(e) \cup \{\vec{0}\}.$$

More generally, we say h is a **virtual configuration** of τ if for each two-ended edge e

$$h(\text{head}(e)) - h(\text{tail}(e)) \in \mathbb{R} \cdot d(e).$$

The set of virtual configurations is a linear subspace of the vector space of all maps $V_\Gamma \rightarrow B$. The set of configurations is a closed polyhedral cone in this subspace, cut out by the conditions that the edges be of nonnegative length; we call it the **configuration space or cone of configurations** of the tinkertoy τ . We can use the vector space structure to define the **sum** $h_1 + h_2$ of two (virtual) configurations.

If B is endowed with a lattice, one can speak of **lattice configurations** of the tinkertoy: these are the ones such that the map h takes V_Γ to lattice points in B .

Example 2. *Configurations of polytope tinkertoys* (for cognoscenti of toric varieties only – we neither use nor prove the statements in this example). In the case P a convex lattice polytope such that the edges from each vertex give a \mathbb{Z} -basis of the lattice, there is an associated smooth toric variety, and the polytope tinkertoy is just a way of encoding the (complete) fan of the polytope. The vector space of virtual configurations of the corresponding polytope tinkertoy can be naturally identified with the second equivariant cohomology group of the toric variety [GZ]; the cone of actual configurations is then identified with the equivariant Kähler cone, and the lattice configurations with the equivariant Chern classes of nef line bundles.⁶

We define a **subtinkertoy** $(B, \Delta, d|_{E_\Delta}) \leq (B, \Gamma, d)$ as a tinkertoy living in the same space B , with any subgraph $\Delta \leq \Gamma$, and the same assigned directions d (restricted to the subset E_Δ). We will not have much need for morphisms of tinkertoys, but we do define an **isomorphism** between two tinkertoys $(B, \Gamma_1, d_1), (B, \Gamma_2, d_2)$ in the same space B as a correspondence between the two graphs, intertwining the direction maps d_1, d_2 .

Example 3. *The Gelfand-Cetlin tinkertoy.* Let $B = \mathbb{R}$, $V = \{v_{i,j}\}$ for $1 \leq i \leq j \leq n$, and E consist of two groups of edges $\{e_{i,j}, f_{i,j}\}$, each $1 \leq i \leq j \leq n - 1$. Every edge is assigned the direction \mathbb{R}^+ .

$$\begin{aligned} \text{head}(e_{i,j}) &= v_{i,j+1}, & \text{tail}(e_{i,j}) &= v_{i,j}, \\ \text{head}(f_{i,j}) &= v_{i,j}, & \text{tail}(f_{i,j}) &= v_{i+1,j+1}. \end{aligned}$$

One important subtinkertoy in this consists of the “primary” vertices $\{v_{i,n}\}$ and no edges. The configurations of the Gelfand-Cetlin tinkertoy restricting to a given configuration of the primary vertices form a polytope called the Gelfand-Cetlin polytope. Each high weight of $GL_n(\mathbb{C})$ gives a (weakly decreasing) list of integers,

⁶In order to model line bundles on *noncompact* toric varieties using tinkertoys, we would need a more ornate definition of tinkertoy including higher-dimensional objects than edges, and also a more ornate definition of configuration, assigning affine subspaces to all objects in the tinkertoy (not just to the vertices). These additional complications only serve to obscure the simplicity of the tinkertoys actually used in this paper.

which we take as a lattice configuration of the primary vertices. The lattice points in the Gel'fand-Cetlin polytope are called Gel'fand-Cetlin patterns, and they count the dimension of the corresponding irreducible representation of $GL_n(\mathbb{C})$. Note that not every configuration of the primary vertices can be extended to a configuration of the whole Gel'fand-Cetlin tinkertoy – for this to be possible, the coordinates of the primary vertices *must* be weakly decreasing.

Return now to the general case. If the two-ended edges e of a tinkertoy τ are all of *positive* length in a configuration h , or equivalently

$$h(\text{head}(e)) - h(\text{tail}(e)) \in d(e),$$

we call the configuration h **nondegenerate**.⁷ Otherwise we say that h is a **degenerate configuration**, each edge e with $h(\text{head}(e)) = h(\text{tail}(e))$ is a **degenerate edge** of h , and each vertex attached to a degenerate edge is a **degenerate vertex** of h . Note that not every tinkertoy *has* a nondegenerate configuration – for example, make a tinkertoy with vertices x, y and two edges e, f from x to y with different given directions $d(e) \neq d(f)$.

The following proposition is of a type standard in convex geometry:

Proposition 1. *If a tinkertoy τ has a nondegenerate configuration, then the nondegenerate configurations form the interior of the cone of configurations, and every configuration is a limit of nondegenerate ones. If τ doesn't have any nondegenerate configurations, there is some edge e that is degenerate in every configuration of τ .*

It is worth noting that if h is a degenerate configuration of a tinkertoy τ , one can associate a smaller tinkertoy $\bar{\tau}$ in which all the degenerate edges of h have been removed, and any two vertices connected by a series of degenerate edges have been identified; the configuration h then descends to a *nondegenerate* configuration of this smaller tinkertoy. In this way each face of the cone of configurations can be identified with the full cone of configurations of a smaller tinkertoy.

2.2. The relevant vector space B for this paper's tinkertoys. From here on out, all our tinkertoys are going to live in the same space $\mathbb{R}_{\Sigma=0}^3 := \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 0\}$, a plane containing the triangular lattice $\mathbb{Z}_{\Sigma=0}^3$. This plane has three **coordinate directions** $(0, -1, 1)$, $(1, 0, -1)$, $(-1, 1, 0)$, and each direction $d(e)$ will be one of these.

In particular, as one traverses the interval assigned to an edge by a configuration, one coordinate remains constant while the other two trade off, maintaining zero sum. We will call this the **constant coordinate** of the edge in the configuration.

Example 4. *The GL_2 honeycomb tinkertoy, in Figure 1.* (This case is too small to see why these are named “honeycombs”.) This tinkertoy has one vertex which is the tail of three edges in the three coordinate directions, three vertices that are each the head of three such edges, for a total of four vertices and nine edges (six of which have no tails and are thus semi-infinite).

The constant coordinates on the six semi-infinite edges determine the configuration, and looking at the central vertex, we see that their sum must be zero. This equality is only sufficient for the existence of a *virtual* configuration: any

⁷One last toric variety remark: in the context of the configurations of polytope tinkertoys, nondegenerate lattice configurations correspond to ample line bundles on the corresponding projective toric variety.




FIGURE 1. A configuration of the GL_2 honeycomb tinkertoy, the vertices labeled with their coordinates in $\mathbb{R}^3_{\sum=0}$. The lengths of the two-ended edges are given at right, in terms of the separations λ, μ, ν .

actual configuration will also satisfy the triangle inequalities on the separations $\lambda = \lambda_1 - \lambda_2, \mu = \mu_1 - \mu_2, \nu = \nu_1 - \nu_2$ between the pairs of semi-infinite edges going in a coordinate direction.

There is a concise way to describe the set of integral coordinates $\lambda_1, \lambda_2, \mu_1, \mu_2, \nu_1, \nu_2$ that arise in configurations of the GL_2 honeycomb tinkertoy: they are exactly those such that the tensor product $V_{(\lambda_1, \lambda_2)} \otimes V_{(\mu_1, \mu_2)} \otimes V_{(\nu_1, \nu_2)}$ of the corresponding representations of GL_2 contains an invariant vector. (Proof sketch: the requirement that the sum be zero is equivalent to asking that the center of GL_2 act trivially on the tensor product. Then the triangle inequalities are familiar from SL_2 theory.) In these cases, the configuration is unique, and so too is the invariant vector (up to scale).

Much of the rest of this section is about generalizing this example to general GL_n , which quite amazingly can also be performed in the plane $\mathbb{R}^3_{\sum=0}$.

2.3. The infinite honeycomb tinkertoy, and GL_n honeycomb tinkertoys. As promised, $B = \mathbb{R}^3_{\sum=0}$. Let V be the set of points⁸

$$V := \{(i, j, k) \in \mathbb{Z}^3_{\sum=0} : 3 \text{ doesn't divide } 2i + j\}.$$

For each vertex $(i, j, k) \in V$ such that $2i + j \equiv 2 \pmod 3$, put on three outwardly directed edges, ending at the vertices $(i - 1, j + 1, k), (i, j - 1, k + 1), (i + 1, j, k - 1)$. These will be the vertices and edges of a directed graph Γ , in which every edge is two-ended, and each vertex has three attached edges, either all in or all out (depending on $2i + j \pmod 3$).

The **infinite honeycomb tinkertoy** is then $(\mathbb{R}^3_{\sum=0}, \Gamma, d)$, where the direction $d(e)$ of an edge is its direction $(head(e) - tail(e))/\mathbb{R}^+$, and the inclusion map $V \hookrightarrow \mathbb{R}^3_{\sum=0}$ defines a nondegenerate configuration of this tinkertoy; see Figure 2.

⁸This is sl_3 's weight lattice minus its root lattice. Presumably there is a deep meaning to this – perhaps relating to the A_2 web diagrams in [Ku] – but we did not uncover it.




FIGURE 2. A small region in the standard configuration of the infinite honeycomb tinkertoy. The six adornments on the boundary show how the coordinates change as one moves in that direction.




FIGURE 3. The triangle in the infinite honeycomb tinkertoy containing the GL_4 honeycomb tinkertoy, and that tinkertoy on its own.

The GL_n **honeycomb tinkertoy** τ_n is the subtinkertoy of the infinite honeycomb tinkertoy whose vertices are the $(i, j, k) \in V$ contained in the triangle $j + 3n \geq i \geq k \geq j$ (automatically in the interior), and all their attached edges. This tinkertoy has $3n$ tailless edges; we will call these the **boundary** edges of this tinkertoy. The $n = 4$ example can be seen in Figure 3.

There is a more general notion of honeycomb tinkertoy that we defer until section 3.

We call a configuration h of τ_n a **honeycomb** or τ_n -**honeycomb**. It is important to distinguish the ontological levels here – the GL_n honeycomb *tinkertoy* is an abstract graph with some labeling by directions, whereas a *honeycomb* is the additional data of an actual configuration of that tinkertoy in $\mathbb{R}^3_{\sum=0}$.

Given a τ_n -honeycomb h , we can read off the constant coordinates on the $3n$ semi-infinite edges starting from the southwest and proceeding clockwise.⁹ Denote these $\lambda_1, \dots, \lambda_n, \mu_1, \dots, \mu_n, \nu_1, \dots, \nu_n$, as in Figure 4; these are the **boundary conditions** of the honeycomb.




FIGURE 4. The constant coordinates on the boundary edges of a τ_5 -honeycomb. (The stars are the nonconstant coordinates.)

Let $\text{HONEY}(\tau_n)$ denote the cone of τ_n -honeycombs, and $\text{BDRY}(\tau_n)$ the cone of possible boundary conditions (λ, μ, ν) of τ_n -honeycombs. That is to say, $\text{BDRY}(\tau_n)$ is the image in $(\mathbb{R}^n)^3$ of the map $h \mapsto$ its boundary conditions λ, μ, ν .

Our purpose in introducing honeycombs is to calculate Littlewood-Richardson coefficients, the dimensions $\dim(V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})}$. We do this by linearly relating τ_n -honeycombs to Berenstein-Zelevinsky patterns¹⁰ in an appendix, where we establish the \mathbb{Z} -linear equivalence of the Berenstein-Zelevinsky cone with the space $\text{HONEY}(\tau_n)$.

That equivalence has the following consequence:

Theorem (from Appendix 1). *Let λ, μ, ν be a triple of dominant weights of $GL_n(\mathbb{C})$, and τ_n the GL_n honeycomb tinkertoy. Then the number of lattice τ_n -honeycombs whose semi-infinite edges have constant coordinates $\lambda_1, \dots, \lambda_n, \mu_1, \dots, \mu_n, \nu_1, \dots, \nu_n$ as in Figure 4 is the Littlewood-Richardson coefficient $\dim(V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})}$.*

This generalizes the GL_2 case we did before as Example 4.

⁹If our notion of “configuration” of a tinkertoy explicitly assigned lines in B to edges in E_Γ , not just to those with a head or tail, we could regard this as the restriction of a configuration of the GL_n honeycomb tinkertoy to the subtinkertoy consisting of the boundary edges and no vertices.

¹⁰Gleizer and Postnikov have recently given a way of relating honeycomb configurations to Berenstein-Zelevinsky patterns that is very different from ours [GP].




FIGURE 5. The honeycombs computing the tensor square of $GL_3(\mathbb{C})$'s adjoint representation. N.B. the semi-infinite edges are labeled with their constant coordinates – usually we will label them with their multiplicities.

Example 5. In Figure 5 we calculate the tensor square of the adjoint representation of GL_3 . The corresponding Littlewood-Richardson rule calculation, throwing away partitions with more than three rows, gives

$$V_{(2,1,0)} \otimes V_{(2,1,0)} = V_{(4,2,0)} \oplus V_{(3,2,1)}^{\oplus 2} \oplus V_{(4,1,1)} \oplus V_{(3,3,0)} \oplus V_{(2,2,2)}.$$

(Recall that to turn the S_3 -symmetric honeycomb formulation back into a tensor product decomposition, one must reverse and negate the weight considered the “output”.)

One elementary consequence of this theorem is that the sum of the constant coordinates on the boundary edges is zero. (Proof: by linearity and continuity, it is enough to check on lattice τ_n -honeycombs. On the representation theory side, the sum of the constant coordinates gives the weight of the action of the center of $GL_n(\mathbb{C})$, which must be trivial for there to be any invariant vectors. QED) We now set up a more direct proof by a sort of Green’s theorem argument, proving some other results in tandem.

For an edge e in the GL_n honeycomb tinkertoy τ_n , let the closed **interval** $I_{h,e}$ be defined by

$$I_{h,e} := \left(h(\text{head}(e)) - \mathbb{R}_{\geq 0} \cdot d(e) \right) \cap \left(h(\text{tail}(e)) + \mathbb{R}_{\geq 0} \cdot d(e) \right)$$

$\bullet \longrightarrow \qquad \qquad \qquad \longleftarrow \bullet$

where if $\text{head}(e)$ or $\text{tail}(e)$ is undefined the corresponding term is omitted. (Note that τ_n has no zero-ended edges, so this intersection is never over the empty set.) Say that a curve H in $\mathbb{R}_{\sum=0}^3$ **intersects the configuration h transversely** if H contains no points $h(v)$, and intersects each interval $I_{h,e}$ transversely or not at all.

Lemma 1. *Let τ_n be the GL_n honeycomb tinkertoy,¹¹ h a τ_n -honeycomb, and γ a piecewise-linear Jordan curve in $\mathbb{R}_{\sum=0}^3$ intersecting h transversely. For each edge e let γ_e be the number of times $I_{h,e}$ pokes through γ from the inside to the outside, minus the number from the outside to the inside (the total¹² will be 1, 0, or -1).*

1. *The sum over $e \in E$ of the unit vectors in the directions $d(e)$, weighted by γ_e , is the zero vector.*
2. *The sum over $e \in E$ of the constant coordinates on e , weighted by γ_e , is zero.*

Proof. Since τ_n has its standard configuration, which is nondegenerate, by Proposition 1 the nondegenerate τ_n -honeycombs are open dense in the cone of all τ_n -honeycombs. The space of τ -honeycombs intersecting γ transversely is open in the

¹¹This lemma applies word-for-word to the more general honeycomb tinkertoys defined in the next section.

¹²This is also equal to the “number of heads of e landing inside γ minus the number of tails” – a perhaps misleading phrase, since each number is only ever zero or one!




FIGURE 6. Replacing a single Jordan curve by two, with the dashed line as the shortcut.


space of all τ -honeycombs, and each of the functionals above is obviously piecewise linear (and continuous) on it. Therefore the nondegenerate τ -honeycombs are open dense in the ones intersecting γ transversely, and by continuity it suffices to check the lemma for them.

If γ encloses one (or no) vertices the statement is easily checked. Otherwise we can connect two points on γ by a path within the interior intersecting h transversely and separating the vertices into two smaller groups (see Figure 6). This shortcut gives us two new Jordan curves, γ_1 and γ_2 . One checks that each of the above functionals satisfies $f(\gamma) = f(\gamma_1) + f(\gamma_2)$. By induction the two terms on the right-hand side are zero, and therefore the left is also. \square

In particular, if we take our Jordan curve to be (a PL approximation to) a very big circle, we recover the previous result that the sum of the constant coordinates on the boundary edges is zero.

2.4. Eliding simple degeneracies. Recall from above that we call a tinkertoy configuration h degenerate if some edge has length zero. This can be regarded as a configuration of a simpler tinkertoy $\bar{\tau}$, in which the two vertices collapsed together have been identified (and the edge removed). The configuration map $h : V_{\bar{\tau}} \rightarrow B$ descends to give a configuration \bar{h} of $\bar{\tau}$. In this way the faces of a configuration cone can be identified with configuration cones of simpler tinkertoys.

The case of interest to us is when a single edge of a honeycomb tinkertoy degenerates to a point – or more generally, when no two degenerate edges share a vertex. In this very special case there is an even simpler tinkertoy to consider, where these five edges E, F, G, H, I and two vertices x, y are not replaced by four edges E, F, H, I and one vertex $x = y$, but two edges $E = I, F = H$ and no vertices at all. (In particular, though E 's head/ I 's tail x is removed, the identified edge $E = I$ gets the head r and tail p ; similarly $F = H$ gets the head s and tail q .)



We will call this modification of a tinkertoy **eliding** the edge G , or the vertices x, y , or just eliding the singularity. A tinkertoy created by eliding a number of

edges in a honeycomb tinkertoy we call a **post-elision tinkertoy**. Note that there is no analogue of this for a degenerate edge in a general tinkertoy – it is crucial that $d(E) = d(I), d(F) = d(H)$ so that the d map is well defined on the post-elision tinkertoy. Note also that worse degenerations, as occur in the tensor product calculation in Figure 5, do not usually allow the vertex to be removed.

We will say a honeycomb h has only **simple degeneracies** if no two degenerate edges meet in a vertex. In this case we will typically elide the degenerate edges in the sense of the paragraph above.

It is not readily apparent what the degrees of freedom of a tinkertoy are. However, for post-elision tinkertoys one can say something useful.

Lemma 2. *Let τ be the GL_n honeycomb tinkertoy.¹³ Let h be a τ -honeycomb, some of whose degeneracies are simple, and let $\bar{\tau}$ be a post-elision tinkertoy obtained by eliding some of h 's simple degeneracies, so h descends to a configuration \bar{h} of $\bar{\tau}$. Let γ be a loop (undirected) in the underlying graph of $\bar{\tau}$ containing only nondegenerate vertices of $\bar{\tau}$ (necessarily trivalent). Then there is a one-dimensional family of configurations of $\bar{\tau}$, starting from \bar{h} , in which one moves only the vertices in γ .*

Proof. We can assume that the loop doesn't repeat vertices; if it does, it will have subloops that do not (in which case we will move a proper subset of γ 's vertices).

Orient the loop, and label the vertices with signs based on whether the loop turns left or right at the vertex, as in Figure 7. Because of the angles, if all the left-turn vertices move so as to shrink their nonloop edges by a fixed length ϵ , whereas the right-turn vertices move so as to extend their outgoing edge by the same ϵ , the angles remain unchanged – i.e. we have a new configuration. \square




FIGURE 7. Successive vertices turning the same direction, and in opposite directions, and where they could move (dashed). The signs indicate the change in length of the nonloop edge at the vertex.

Note that no loop can go through semi-infinite edges. So in this previous lemma we're only studying degrees of freedom which leave the semi-infinite edges in place. Also, because we can orient the loop either way, the loop can breathe both in and out.

¹³Again, this lemma extends word-for-word to the general honeycomb tinkertoys defined later.

3. THE DIAGRAM AND DEGENERACY GRAPH OF A HONEYCOMB,
AND RECONSTRUCTING A HONEYCOMB FROM ITS DIAGRAM

3.1. **Honeycomb tinkertoys and honeycomb diagrams.** For h a configuration of a tinkertoy $\tau = (B, \Gamma, d)$, define the **diagram** m_h of the configuration h to be a measure on B , the sum

$$m_h := \sum_{e \in E_\Gamma} \text{Lebesgue measure on the interval } I_{h,e}.$$

This is a little more information than the set-theoretic union $\bigcup_{e \in E_\Gamma} I_{h,e}$, in that it remembers multiplicities when edges are directly overlaid. Note that we can recover the union from m_h , as its support $\text{supp } m_h$.

In this section we will prove that a τ_n -honeycomb is reconstructible from its diagram. There is a stronger statement – that every measure on $\mathbb{R}^3_{\sum=0}$ that looks enough like the diagram of a honeycomb is indeed the diagram of a unique honeycomb (up to a trivial equivalence) – but it requires a more general definition of honeycomb tinkertoy than the GL_n honeycomb tinkertoys we have met so far.

Define a **hexagon** in the infinite honeycomb tinkertoy as the six vertices around a hole, i.e. a 6-tuple of vertices $(i - 1, j + 1, k), (i, j - 1, k + 1), (i + 1, j, k - 1), (i + 1, j - 1, k), (i, j + 1, k - 1), (i - 1, j, k + 1)$ where 3 divides $2i + j$.

Define a **honeycomb tinkertoy** τ as a subtinkertoy of the infinite honeycomb tinkertoy satisfying five conditions:

1. τ is finite;
2. (the underlying graph of) τ is connected;
3. each vertex in τ has all three of its edges (which may now be one-ended);
4. τ contains a vertex (it’s not just a single no-ended edge);
5. if four vertices of a hexagon are in τ , all six are.

(It is slightly unfortunate to rule out the infinite honeycomb tinkertoy itself, but it would be more unfortunate to have to say “finite honeycomb tinkertoy” throughout the paper.) We will call the configuration of τ restricted from the defining configuration of the infinite honeycomb tinkertoy the **standard configuration** of τ .

A honeycomb tinkertoy has a number of semi-infinite edges in each of the three coordinate directions and their negatives. Call this ordered 6-tuple, counted clockwise from the North, the **type** of the honeycomb tinkertoy. One fact we prove later, in Lemma 6, is that any two honeycomb tinkertoys of the same type are isomorphic – that the honeycomb tinkertoys presented in Figure 8 essentially capture all the types. (In fact they are better than ‘isomorphic’; they differ only by translation within the infinite honeycomb tinkertoy.) We give the more axiomatic definition above to make it easy to check, in Lemma 7, that a certain subtinkertoy is itself a honeycomb tinkertoy.




FIGURE 8. The standard configurations of some honeycombs. The boundary edges are semi-infinite.

If τ is a honeycomb tinkertoy, we define a τ -**honeycomb** h to be a configuration of τ , and will speak of h 's type (meaning the type of τ). Just as in the case of the GL_n honeycomb tinkertoy τ_n , we use $\text{HONEY}(\tau)$ to denote the cone of τ -honeycombs, and $\text{BDRY}(\tau)$ to denote the cone of possible constant coordinates on the set of one-ended edges of τ .

It is quite easy to describe the local structure of the diagram of a (perhaps degenerate) honeycomb h . In the pictures to follow of honeycomb diagrams we label edges with their multiplicities (multiple of Lebesgue measure on the line).

Lemma 3. *Let m_h be the diagram of a honeycomb h . Each point $b \in \mathbb{R}^3_{\Sigma=0}$ has a neighborhood in which m_h satisfies one of the following:*

1. $m_h = 0$.
2. m_h is equal to a natural times Lebesgue measure on a coordinate line through b .
3. m_h matches one of the vertices in Figure 9, up to rotation (here the edge multiplicities are naturals).

There are only finitely many points in the third class, the **vertices** of the honeycomb diagram m_h . If one thinks of each of the edges meeting a vertex as pulling on the vertex with a tension equal to its multiplicity, these are exactly the ways for the vertex to experience zero total force.




FIGURE 9. The **Y**, crossing, rake, 5-valent, and 6-valent vertices.

Proof. That the pictures in Figure 9 are the only possibilities, and the finiteness of the number of points in the third class, each follow from the finiteness of the edge set of a honeycomb tinkertoy (so the neighborhood can be shrunk to avoid the $I_{h,e}$ not actually meeting the point b), and the directions $d(e)$ being multiples of 60° from the North. It remains to be sure the multiplicities are constrained as claimed.

By Lemma 1, the sum of the unit outgoing edges of a vertex (weighted by their multiplicities) must be zero. So if both a direction and its negative appear with positive multiplicity, we can subtract one from each and continue. Eventually we must get to a vertex of the first type, **Y**, or nothing at all. \square

Each of the points b of the third class in the above lemma we will call a **vertex of the diagram** m_h .

To see that each of these vertex types actually occurs, start with the honeycombs in Figure 8 (or larger versions with more hexagons) and degenerate all the two-ended edges to points. As we will show in this section, collapsing the honeycombs in Figure 8 is essentially the *only* way to produce the singular vertices in Lemma 9.

3.2. The dual graph $D(\tau)$ of a honeycomb tinkertoy τ . Call the lattice $\{(i, j, k) \in \mathbb{Z}^3_{\Sigma=0} : 3 \text{ divides } 2i + j\}$ the **root lattice** (in that it is the root lattice of \mathfrak{sl}_3), and define the **root lattice triangle** around a vertex (i, j, k) of the infinite honeycomb tinkertoy to be the three points of the root lattice at L^1 -distance




FIGURE 10. A honeycomb tinkertoy τ of type $(7, 0, 4, 5, 2, 2)$ in solid lines and its corresponding $D(\tau)$ in dashed lines, shown superimposed in $\mathbb{R}_{\Sigma=0}^3$.

2 from (i, j, k) (these are the three closest points). We will need the small triangular graph made from a root lattice triangle, and also the region enclosed.

Fix a honeycomb tinkertoy τ . Define $D(\tau)$, the **dual graph of τ** , to be the union of the root lattice triangles around the vertices in τ – this has one vertex in each region in the standard configuration of τ (including the unbounded ones), with an edge connecting two $D(\tau)$ -vertices if the corresponding τ -regions share an edge (see Figure 10 for an example). (In this way edges in $D(\tau)$ correspond to perpendicular edges in τ .) In particular $D(\tau)$ is naturally embedded in $\mathbb{R}_{\Sigma=0}^3$ – it has more structure than just the abstract dual graph.

Lemma 4. *Let τ be a honeycomb tinkertoy. Then the region bounded by its dual graph $D(\tau)$ is convex, and is thus characterized (up to translation) by its 6-tuple of edge-lengths. Conversely, every convex union of root lattice triangles arises as a $D(\tau)$.*

Proof. Each vertex in τ gives us three vertices in $D(\tau)$ (by either adding 1 to, or subtracting 1 from, each of the three coordinates), and thus a small triangle. Two connected vertices in τ share two of these three vertices, so their corresponding triangles in $D(\tau)$ intersect in an edge (and not just a vertex).

Since τ is by assumption connected, any two of these triangles are connected by a chain of triangles sharing common edges. This shows that $D(\tau)$ bounds a single region, not a disconnected set, nor two regions intersecting in only a vertex.

It remains to prove this region is convex. If not, it has an internal angle of more than 180° going around some boundary vertex. This would mean four successive triangles out of the same vertex are in $D(\tau)$. On the τ side, that means four successive vertices around a hexagon are in τ . But that forces the whole hexagon to be in τ . So the vertex was not actually on the boundary, a contradiction. This establishes the convexity of $D(\tau)$.

The converse is simple: given a convex union D of root lattice triangles which we wish to realize as a dual graph $D(\tau)$, take τ to be the subtinkertoy of the infinite honeycomb tinkertoy lying within D (those vertices, and all their edges). This is easily seen to be a honeycomb tinkertoy whose dual graph $D(\tau)$ is D . \square

So the external angles are restricted to 0° , 60° , and 120° . To rotate once, the total of the external angles must be 360° , so there are five types, depending on the number and ordering of the two kinds of angles; see Figure 11.




FIGURE 11. The possible shapes of the region bounded by a dual graph $D(\tau)$.

Not any 6-tuple of edge-lengths will do; the boundary of $D(\tau)$ must be a closed curve. This gives two linear conditions that exactly match the zero-tension property of Lemma 3. So we've essentially classified the possible $D(\tau)$.

To get an equally good hold of τ , we need a lemma saying we can reconstruct τ from $D(\tau)$. This will follow from a study of the length-minimizing paths connecting two points in a honeycomb tinkertoy, which we dub **geodesics**. (Note that there are typically a great many paths with this minimum length.)

The following lemma says that τ is “geodesically convex” inside the infinite honeycomb tinkertoy.

Lemma 5. *Let τ be a honeycomb tinkertoy, A, B two vertices of τ , and γ a geodesic between A and B in the infinite honeycomb tinkertoy. Then the vertices of γ are in τ .*

Proof. We induct on the length of γ , assuming that the lemma is proven for all A and B with geodesics shorter than γ .

Since τ is connected, A and B are connected under *some* path δ in τ . Let Ω be the collection of hexagons enclosed by the concatenation $\gamma + \delta$. We can assume the cardinality of Ω is minimal among all possible paths δ in τ that connect A and B . If Ω is empty, we are done, so suppose for contradiction that Ω is nonempty.

If $\gamma + \delta$ is not a Jordan curve, we can break it into smaller pieces and use the induction and minimality hypotheses. Hence we may assume $\gamma + \delta$ is Jordan.

The curve δ cannot contain three consecutive edges of a hexagon in Ω ; if it did, then by property 4 required of honeycomb tinkertoys, all the vertices of this hexagon would be in τ . Then we could “flip” δ to go around the other side of the hexagon, which would remove that hexagon from Ω and contradict the minimality assumption. Thus we may assume that δ does not contain three consecutive edges of any hexagon in Ω .

To finish the contradiction we shall invoke

Sublemma 1. *Suppose $\gamma + \delta$ is a Jordan curve whose interior Ω is a nonempty collection of hexagons. Suppose further that δ does not contain three consecutive edges of a hexagon in Ω . Then γ is longer than δ .*

Proof. Suppose for contradiction that there was a counterexample $\gamma + \delta$ to this sublemma. We may assume that this counterexample has a minimal number of hexagons in Ω . We may assume that Ω contains more than one hexagon, since the sublemma is clearly true otherwise.

From hypothesis, δ does not contain three consecutive edges of any hexagon in Ω . We now claim that γ also does not contain three consecutive edges of any hexagon in Ω . For, if γ did contain three such edges, one could then “flip” these edges across the hexagon; this would preserve the length of γ and therefore contradict

the minimality of Ω . (If the flip operation causes $\gamma + \delta$ to cease being Jordan, eliminate redundant edges and divide into connected components.)

Now traverse $\gamma + \delta$ once in a counterclockwise direction, so that Ω is always to the left. At every vertex one turns 60 degrees in a clockwise or counterclockwise direction. From the above considerations we see that one cannot execute two consecutive counterclockwise turns while staying in the interior of δ , since this would imply that δ contains three consecutive edges of a hexagon in Ω . Similarly one cannot execute two consecutive counterclockwise turns while staying in the interior of γ . Thus, with at most four exceptions, every counterclockwise turn in $\gamma + \delta$ is immediately followed by a clockwise turn. But this contradicts the fact that we must turn 360 degrees counterclockwise as we traverse $\gamma + \delta$. \square

Since γ was assumed to be a geodesic, we have the desired contradiction. \square

Lemma 6. *One can reconstruct a honeycomb tinkertoy τ from its dual graph $D(\tau)$: its vertices are*

$$\{(i, j, k) \in \mathbb{Z}_{\sum=0}^3 : 3 \text{ doesn't divide } 2i + j, \\ \text{the root lattice triangle around } (i, j, k) \text{ is in } D(\tau)\}.$$

The number of semi-infinite edges of τ in a particular direction is equal to the length of the corresponding edge of $D(\tau)$. In particular, honeycomb tinkertoys are characterized (up to translation in the infinite honeycomb tinkertoy) by their type.

Proof. Each point in τ is in the set above, tautologically – a point in τ leads to the three points in $D(\tau)$, which lead back to the same point being in the set above.

Now fix a triangle T in $D(\tau)$; we wish to show that the center of that triangle is necessarily in τ . From a vertex in $D(\tau)$ one can infer that at least one of the six neighboring points in $\{(i, j, k) : i + j + k = 0, 3 \text{ doesn't divide } 2i + j\}$ is in τ . For example, the presence of the dotted (North) vertex in Figure 12 says that up to left-right reflection, there must be a τ vertex in one of the regions labeled 1, 2, 3 or 4.

If there's a τ vertex in region 1, we're done. If there's a τ vertex in region 2 or 3, and another τ vertex producing the existence of the Southeast vertex of $D(\tau)$, we can find (in Figure 13) a geodesic in the infinite honeycomb connecting the two that goes through the center. Then by Lemma 5 about such geodesics, the center is necessarily a vertex of τ .




FIGURE 12. Possibilities for the τ vertex causing the North $D(\tau)$ vertex in this triangle.




FIGURE 13. Geodesics through the center in cases 2,3, and through 2,3 in case 4.




FIGURE 14. A very degenerate honeycomb h , and the corresponding $D(h)$.

The remaining case occurs when the only τ vertex can be found in region 4, in all three rotations of the diagram. But then by connecting two of those τ vertices with geodesics we find τ vertices in regions 2 and 3, reducing to the previous case. \square

3.3. The degeneracy graph $D(h)$ of a honeycomb h . Fix a honeycomb tinkertoy τ , and let $D(\tau)$ be its dual graph. For h a τ -honeycomb, let $D(h)$ be the subgraph of $D(\tau)$ with the same vertex set but an edge between two vertices only if the corresponding (perpendicular) edge of h is nonzero (see Figure 14 for an example). This we will call the **degeneracy graph** of the honeycomb h (and is also embedded in $\mathbb{R}^3_{\Sigma=0}$).

Our goal is to classify the regions in $D(h)$, and establish that they correspond to vertices in the diagram of h , thereby classifying the possible vertices in the diagram and their preimages under h (à la Figure 8).

Lemma 7. *The regions in the degeneracy graph $D(h)$ of a honeycomb h are convex.*

Proof. Fix a region Ω in $D(h)$, and choose a vertex x of τ which lies in the closure of Ω . Then its image $h(x)$ is in the support of the diagram of h .

Consider the subtinkertoy of τ consisting of those vertices in τ which map to $h(x)$ under h , together with their associated edges. Let σ be the connected component¹⁴

¹⁴Actually, we shall see in the next lemma that the tinkertoy necessarily has only one connected component.




FIGURE 15. A straight line from region X to region Y in $D(h)$, and the corresponding path in h .

of this tinkertoy that contains x . Then σ is a honeycomb tinkertoy in its own right: the only nontrivial observation required is that if four vertices of a hexagon map to $h(x)$, then all six vertices must map to $h(x)$.

Chasing down the definitions we see that the region bounded by $D(\sigma)$ is just Ω . The claim then follows from Lemma 4. \square

Lemma 8. *The regions in the degeneracy graph $D(h)$ of a honeycomb h correspond to the vertices in the diagram m_h of h .*

Proof. One direction is clear: if two vertices in τ give root lattice triangles in the same region of $D(h)$, those vertices are connected by a series of degenerate edges in h , and therefore have collapsed to the same vertex of the diagram of h . So for each region X in $D(h)$, we can speak of X 's h -vertex. The converse is more difficult; we must show that two distinct regions X and Y in $D(h)$ give vertices in the diagram of h that are physically separated.

One case is easy. If X and Y are adjoining regions, then the existence of the edge in $D(h)$ separating them says that the corresponding edge in h is nonzero. But this is exactly the displacement between X 's h -vertex and Y 's h -vertex, so they are in different places, as desired. For nonadjoining X and Y we will have to add together a bunch of such nonzero displacements and hope to get a nonzero sum.

Let x be a generic point in X , and y a generic point in Y , and let \overline{xy} be the straight-line path connecting them. Then \overline{xy} does not intersect the vertices of $D(h)$ and only intersects the edges of $D(h)$ transversely.

We can now compute the vector in $\mathbb{R}_{\sum=0}^3$ separating the points in h corresponding to the regions X and Y . Each time \overline{xy} crosses a wall – an edge still left in $D(h)$ – there is an associated nonzero displacement, the vector difference of the h -vertices of the regions on the two sides of the wall. Adding up all these displacements we get the total vector difference we seek, the displacement of the vertices corresponding to X and Y . (It is worth emphasizing that this gives a path in h itself, as indicated in Figure 15.) But note now that each individual term has positive dot product with the vector $y - x$, since it is perpendicular to the wall that \overline{xy} has just crossed through (and in the correct direction).

So therefore the whole sum has positive dot product with $y - x$, so is nonzero, and therefore X 's h -vertex and Y 's h -vertex are not in the same place in the diagram. \square

In particular, since the preimage of a vertex of a honeycomb is connected, it is a honeycomb tinkertoy in its own right.

There is a succinct way to sum up the results of this section. Define an **abstract honeycomb diagram**¹⁵ as a measure m on $\mathbb{R}_{\Sigma=0}^3$ such that

1. in a neighborhood of each point $b \in \mathbb{R}_{\Sigma=0}^3$, m is a nonnegative *real* linear combination of the Lebesgue measures on the six coordinate rays out of b , satisfying the zero-tension property of Lemma 3;
2. only finitely many points, which we naturally call the **vertices** of m , use more than two rays in this combination;
3. $\text{supp } m$ is not a collection of parallel lines.

By Lemma 3, the diagram of a honeycomb is an abstract honeycomb diagram. The following theorem states that, up to a trivial ambiguity, every abstract honeycomb diagram with integral edge-multiplicities is the diagram of a unique honeycomb.

Theorem 1. *Let m be an abstract honeycomb diagram with integral edge-multiplicities. Then m is the diagram of a configuration h of a honeycomb tinkertoy τ , where τ is uniquely determined by m up to unique isomorphism; and given τ the honeycomb h is unique.*

Proof. We break the diagram-to-honeycomb reconstruction into two steps: from the diagram to the degeneracy graph, then from the degeneracy graph to the honeycomb.

Finding the degeneracy graph $D(h)$. We first need to determine the dual graph $D(\tau)$ containing $D(h)$. Since each of the vertices in m satisfies the zero-tension property of Lemma 3, the whole diagram does, by the same Green's theorem argument as in Lemma 1. Correspondingly, there does exist a convex lattice region whose edge-lengths are the numbers of semi-infinite edges, unique up to translation.

By Lemma 8, the vertices in m are supposed to correspond to the regions in $D(h)$, and we can determine the shapes of those regions from Lemma 6. It remains to fit them all together into $D(\tau)$.

For each vertex p in m , pick a path in $\text{supp } m$ whose last vertex connects to a semi-infinite edge. Each vertex along the path corresponds to a region in $D(h)$ whose shape we can determine from the vertex and Lemma 6. These regions glue together along edges perpendicular to the steps in the path. We can determine where the region corresponding to the last vertex in the path sits in $D(h)$: we know it's on a boundary edge of $D(h)$ (the boundary corresponding to the direction of the semi-infinite edge), and we know where it sits on that edge, by counting how many semi-infinite edges (with multiplicity) going in that direction are to the right and left of this bunch. Having done that, by gluing the other regions to it we have determined where they all sit, including that of the original vertex.

The only worry then is that different paths to the boundary may suggest different places to locate p 's region inside $D(h)$. One checks that there is no monodromy in going around a small loop in the embedded graph $\text{supp } m$ – this is because the hexagon-closure condition of a small path is the same as the zero-tension condition of Lemma 3 – and therefore none in any loop.

¹⁵This concept matches the *web functions* of [GP]; we prefer to avoid this terminology, though, for fear of confusion with the intriguingly similar ‘webs’ of [Ku]. A. Postnikov of [GP] has informed us that he also knew Theorem 1.

Finding the honeycomb tinkertoy τ and the honeycomb h . From $D(\tau)$, using Lemma 6 we can construct τ uniquely. Had $D(\tau)$ been chosen differently, τ would only change by translation within the infinite honeycomb tinkertoy.

We've determined $D(h)$'s breakup into regions, and which vertex in p corresponds to which region in $D(h)$. But this determines the honeycomb h – for each vertex v of p , take the vertices of τ in the interior of the $D(h)$ -region corresponding to v and map them to v under h . \square

The degeneracy graph $D(h)$ is thus a way of recording only what we might call the combinatorial information about a honeycomb h , not the actual positions. (In fact there is a tighter connection: see Appendix 2.) The vertices of one correspond to regions in the other, and the length of an edge of $D(h)$ is equal to the multiplicity of the corresponding (perpendicular) edge in h .¹⁶

This theorem allows us to work pretty interchangeably with honeycombs vs. honeycomb diagrams. In particular we can strengthen the result of Lemma 2, which only gave us a family of configurations of post-elision tinkertoys $\bar{\tau}$ made from honeycomb tinkertoys τ , to actually give us configurations of τ itself.

Corollary. *Let h be a τ -honeycomb, and let $\bar{\tau}$ be a tinkertoy obtained by eliding some of h 's degenerate edges, so h descends to a configuration \bar{h} of $\bar{\tau}$. Let γ be a loop (undirected) in the underlying graph of $\bar{\tau}$ passing only through nondegenerate vertices of \bar{h} . Then there is a one-dimensional family of configurations of τ , starting from h , in which one moves only the vertices in γ .*

Proof. Apply Lemma 2 to get a family of configurations of $\bar{\tau}$, and apply Theorem 1 to their diagrams; this produces configurations of τ . \square

This is most interesting when breathing the loop in and out causes a vertex of the post-elision tinkertoy $\bar{\tau}$ to move across an edge, as in Figure 16, something not possible for the honest honeycomb tinkertoy τ . In this case, the application of Theorem 1 to the 1-dimensional family of $\bar{\tau}$ configuration diagrams produces a *piecewise-linear* 1-dimensional family of τ -honeycombs that bends around the cone $\text{HONEY}(\tau)$.




FIGURE 16. Two members of a bent family of GL_4 honeycombs and their degeneracy graphs, resulting from breathing a hexagon in their common post-elision tinkertoy. Since their degeneracy graphs are different, the honeycombs lie on different faces of the cone of honeycombs.

¹⁶If the edges of $D(h)$ are given formal multiplicities equal to the length of the corresponding edge in h , the graph $D(h)$ becomes a honeycomb diagram itself (ignoring here the problem of dealing with the semi-infinite edges). This is essentially the duality in [GP] on BZ triangles; see also the remarks at the end of the next section.

4. OVERLAYING HONEYCOMBS, AND THE PRV CONJECTURE

The reconstruction theorem of the last section lets us define a remarkable operation on honeycombs. (This section is not used elsewhere in this first paper.)

Corollary (to Theorem 1). *Let h and h' be two honeycombs (perhaps of different types). Then (up to translation in the infinite honeycomb tinkertoy) there exists uniquely a honeycomb whose diagram is the sum of the diagrams of h and h' (as measures; one adds multiplicities when the edges of h and h' lie fully on top of one another).*

Proof. Let m be the sum as a measure of the diagrams of h and h' . One checks straightforwardly that m is an abstract honeycomb diagram with integral edge-multiplicities, except for connectedness.

If $\text{supp } m$ is not connected, there is a connected component C of $\mathbb{R}_{\Sigma=0}^3 \setminus \text{supp } m$ bounded by two different components of $\text{supp } m$. The region C is not convex and bounded, or else its boundary would have just one (polygonal) component. If C is not convex it must be nonconvex at one of its vertices, violating Lemma 3's zero-tension condition on m 's vertices, a contradiction.

So C is convex and unbounded, hence its limit points in the circle at infinity $\text{Rays}(\mathbb{R}_{\Sigma=0}^3)$ must either be an interval or two opposite points. If this is an interval, C 's boundary is connected, a contradiction. So C contains a line, and its boundary in $\text{supp } m$ must be two parallel lines. By the zero-tension condition on honeycomb vertices, there can be no vertices on these lines. Removing them, and repeating the argument, we find that h and h' are each unions of parallel lines, contrary to assumption.

So we have the assumed connectedness, and m is an abstract honeycomb diagram. Then by Theorem 1 it is the diagram of an essentially unique honeycomb. \square

The naturality of this *piecewise-linear* operation on honeycombs is, to our minds, one of the principal advantages over the BZ formulations, and will be central in the next paper [Hon2]. (In this paper the overlay notion is only used in a sort of local way – the elision operation on simple degeneracies.)

Application: the weak PRV conjecture for $GL_n(\mathbb{C})$. The so-called weak PRV conjecture (now proven in general [KMP]) states that if $w\lambda + v\mu$ is in the positive Weyl chamber for some Weyl group elements w, v , then $V_{w\lambda+v\mu}$ is a constituent of the tensor product $V_\lambda \otimes V_\mu$. We prove this for $GL_n(\mathbb{C})$ as follows. For $i \in \{1, \dots, n\}$, let h_i be a GL_1 -honeycomb, i.e. a single vertex with three semi-infinite edges coming off, whose coordinates are $(\lambda_{w(i)}, \mu_{v(i)}, -\lambda_{w(i)} - \mu_{v(i)})$. Then overlaying all the $\{h_i\}$, we get a lattice GL_n -honeycomb with boundary conditions λ, μ , and $-(w\lambda + v\mu)$. This lattice honeycomb is a witness to this instance of the weak PRV conjecture (see Figure 17 for an example).

The cost of working in the honeycomb model is that the linear structure – the fact that one can add two BZ patterns of the same size and get another – is geometrically a slightly more mysterious operation on honeycombs. As we will see in a later paper in this series, these two operations “add” and “overlay” are intertwined by the duality operation on honeycombs from [GP], which should be seen as a tropical version of the Fourier transform relating “times” and “convolve”. On honeycombs this operation essentially amounts to replacing a honeycomb h with its degeneracy graph $D(h)$, where each edge of $D(h)$ is given a formal “multiplicity” equal to the length of the corresponding edge in h (thus completing the duality of the two




FIGURE 17. An example of a witness to the PRV conjecture.

graphs). We leave the additional details of how to handle the semi-infinite edges to the later paper.

5. THE LARGEST-LIFT MAP $\text{BDRY}(\tau) \rightarrow \text{HONEY}(\tau)$

In Figure 18 we see a honeycomb with integral boundary that is not itself integral. It is a little more difficult to see that it is in fact an extremal point on the polytope of honeycombs with this boundary, i.e., the constant coordinates on the interior edges are uniquely determined unless one un-degenerates some zero-length edges. The reader can check this by using the fact that the diagram is in $\mathbb{R}_{\sum=0}^3$ to determine the constant coordinates on all the edges except those in the figure eight; then the fact that the line passes through the node on the figure eight ties down the rest of the coordinates.




FIGURE 18. A nonintegral vertex of a honeycomb polytope, with edges (all multiplicity one) labeled with their constant coordinates.

It will turn out that, in a sense, the 6-valent vertex is to blame for this honeycomb's bad behavior. In this section we develop the machinery to find extremal honeycombs with better vertices (and thus better behavior) than this one.

Let P be a (possibly unbounded) polyhedron in a vector space U , $\pi : U \rightarrow V$ a projection that restricts to a proper map $P \rightarrow Q$, and \vec{w} a generic functional on U . Define the “**largest lift**” map $l : Q \rightarrow P$, taking q to the point in $P \cap \pi^{-1}(q)$ with greatest pairing with \vec{w} . By properness, there is a point maximizing this pairing; by genericity of \vec{w} , this point is unique. So the map is well-defined, and in fact is continuous and piecewise-linear ([Zi, pp. 293-4]). In the case of P a tortoise, Q its shadow on the ground at noon, and \vec{w} measuring the height off the ground, the largest lift of a point in the shadow is the corresponding point on the tortoise's upper shell.

Fix a honeycomb tinkertoy τ . We are interested in this largest-lift map in the case of the projection $\text{HONEY}(\tau) \rightarrow \text{BDRY}(\tau)$, which forgets the location of the vertices and finite edges of a honeycomb, remembering only the constant coordinates on the boundary edges.

Proposition 2. *Let τ be the GL_n honeycomb tinkertoy. Then the map $\text{HONEY}(\tau) \rightarrow \text{BDRY}(\tau)$ is proper.*

Proof. This is guaranteed by the correspondence with the BZ cone (in Appendix 1). \square

In fact this map is improper only if τ has semi-infinite edges in all six directions, but we will not need this fact in this paper.

The functional $w_{\text{perim}} : \text{HONEY}(\tau) \rightarrow \mathbb{R}$ is chosen to be a generically weighted sum of the perimeters of the (possibly degenerate) hexagons in the honeycomb, the weighting having a certain “superharmonicity” property. More exactly, let w assign a real number to each of the regions in the tinkertoy τ (vertices of $D(\tau)$), with the properties that

1. for each unbounded region r on the exterior, $w(r) = 0$;
2. for each hexagon α surrounded by regions α_i , $w(\alpha) > \frac{1}{6} \sum_i w(\alpha_i)$;
3. w is chosen generic, subject to these constraints.

(One nongeneric such w can be defined on any $D(\tau)$ by $w(i, j, k) = -i^2 - j^2 - k^2$. But the set of w is open, so we can perturb this one slightly to get a generic w .)

Then we define the **weighted perimeter** of a τ -honeycomb h as

$$w_{\text{perim}}(h) = \sum_{\text{hexagons } \alpha} w(\alpha) \text{perimeter}(\alpha).$$

Since w_{perim} is defined in terms of the perimeter, it's a linear functional on $\text{HONEY}(\tau)$.

Lemma 9. *Let h be a honeycomb in which some hexagon can inflate (moving the vertices of the loop, as in Lemma 2). Then inflating it increases $w_{\text{perim}}(h)$.*

Proof. Inflating the hexagon by distance ϵ increases its perimeter by 6ϵ , while decreasing that of each of its neighbors by ϵ (see Figure 19). The change in $w_{\text{perim}}(h)$ is

$$6g\epsilon - (a + b + c + d + e + f)\epsilon = (6g - (a + b + c + d + e + f))\epsilon > 0.$$

\square




FIGURE 19. An inflating hexagon, with each region α labeled with its weight $w(\alpha)$.

This lemma is best understood in terms of $\text{HONEY}(\tau)$'s linear structure, which while transparent in the original BZ formulations, or the hive model in Appendix 2, is unfortunately rather opaque in the honeycomb model. For each region R in a honeycomb tinkertoy τ , define the **inflation virtual configuration** \vec{i}_R which places the vertices of R around the origin in $\mathbb{R}_{\Sigma=0}^3$ (by translating the hexagon containing R 's vertices to the hexagon in the infinite honeycomb tinkertoy around the origin), but all other vertices *at* the origin. Then the statement “the hexagon H can inflate in h , but gets stuck at a distance s ” is equivalent to “ $h + \epsilon \vec{i}_H$ is in the cone $\text{HONEY}(\tau)$ for $\epsilon \in [0, s]$, but past s this ray leaves the cone”. And Lemma 9 above is exactly the statement (if we extend w_{perim} linearly to the vector space of virtual configurations of τ) that $w_{\text{perim}}(\vec{i}_R) > 0$.

In a particularly degenerate honeycomb it may be impossible to inflate any one hexagon; only certain combinations may be possible. This next slightly technical lemma shows that certain local changes to a honeycomb, which **molt**¹⁷ the degeneracy, can be obtained by inflating several hexagons simultaneously.

Lemma 10. *The virtual configurations associated to the recipes in Figure 20 for molting a degenerate vertex are in each case the sum of a set of inflation virtual configurations \vec{i}_R associated to inflating a certain collection of regions R . More precisely, each molting recipe inflates equally the completely degenerate hexagons, plus the 4-sided regions on the sides corresponding to thick edges.*

Proof. Since these degenerate vertices involve, by definition, a number of hexagons that have collapsed (to lines and even to points), it is rather difficult to see which hexagons must be simultaneously inflated to make the vertex molt. We will use the linear structure to get around this as follows; first add ϵ times the standard configuration of the honeycomb tinkertoy τ . Now there is no degeneracy and we can point out which regions to inflate. Add the corresponding virtual configurations \vec{i}_R to this configuration. Then subtract ϵ times the standard configuration and see that we do get the molted configuration.

Note in particular that if we mark two adjacent regions for the same amount of inflation, their common edge doesn't move at all. So it is simple to see whether an edge moves under simultaneous inflation, since each edge is on the boundary of exactly two regions; it moves if exactly one of them inflates, away from that one.

¹⁷From Webster's: molt : to shed hair, feathers, shell, horns, or an outer layer periodically : to cast off (an outer covering) periodically; specif : to throw off (the old cuticle).




FIGURE 20. How to “molt” a degenerate vertex.




FIGURE 21. From left to right: a **Y** vertex; that vertex expanded with some regions marked for inflation (in gray); the result of inflation; then de-expanded again.

In each of the following pictures we show

1. the vertex (edges labeled with their multiplicities);
2. the standard configuration of (a small example of) the underlying tinkertoy, with certain regions labeled in gray;
3. the result of inflating those regions some distance;
4. the same result, with the standard configuration subtracted off.

*Molting a **Y** vertex* (Figure 21). We determined in Lemma 6 that **Y** vertices result from the collapse of a $(0, n, 0, n, 0, n)$ -type honeycomb tinkertoy. To perform the sort of molt we want here, we inflate all the regions *except* one corner and the opposite side.

Molting a crossing vertex (Figure 22). In this case we inflate all regions except those on the left and right side. Again, we are using Lemma 6 to know precisely what is hiding in the degenerate vertex, as we will again in the remaining cases.

Molting a rake vertex (Figure 23). In this case we inflate all regions except those on the left, right, and top.

Molting a 5-valent or 6-valent vertex. This is more of the same, so we do not take space for the pictures (see Figure 8 for the standard configurations of the tinkertoys). In both the 5- and 6-valent case, we mark all the hexagons for inflation,




FIGURE 22. Molting a crossing vertex by inflating certain regions.




FIGURE 23. Molting a rake vertex by inflating certain regions.

and in the 5-valent case we also mark for inflation the 4-sided unbounded regions on the side with m external edges. \square

Note that we make no statement here about the change in the weighted perimeter, since (except in the 6-valent vertex case) some of the regions we are inflating are unbounded, and Lemma 9 does not pertain. In the theorem to follow we will only be inflating hexagons.

For τ a honeycomb tinkertoy, call $\beta \in \text{BDRY}(\tau)$ a set of **regular** boundary conditions if no two semi-infinite edges of τ in the same direction are assigned the same constant coordinate. This terminology is taken from the GL_n -honeycomb case, where such boundary conditions correspond to triples of regular dominant weights (i.e. in the interior of the positive Weyl chamber).

The main result in this paper is the following.

Theorem 2. *Let τ be a honeycomb tinkertoy such that the map $\text{HONEY}(\tau) \rightarrow \text{BDRY}(\tau)$ is proper, as in Proposition 2. Let β be a regular point in $\text{BDRY}(\tau)$, and $l(\beta)$ the largest-lift honeycomb lying over it (relative to a generic choice w). Then $l(\beta)$ has only simple degeneracies, and if one elides them, the graph underlying the resulting tinkertoy is acyclic.*

Proof. Since the map is assumed proper, the concept of a “largest lift” makes sense, and we can go on to study its properties. In order, we will show

1. a largest lift never has 6-valent vertices;
2. a largest lift of a regular point has no edges of multiplicity > 1 (and therefore no rakes or 5-valent vertices);
3. a simply degenerate largest lift has no cycles.




FIGURE 24. Replacing a multiplicity- m path with a multiplicity- $(m - 2)$ path while molting a multiplicity-1 skin.

Item #1 is immediate from Lemma 10: if $l(\beta)$ has a 6-valent vertex, we can molt it by inflating a certain set of hexagons. But by Lemma 9 that increases the weighted perimeter. So $l(\beta)$ was not a largest lift, contrary to assumption.

Example 6. In the calculation of the tensor square of $GL_3(\mathbb{C})$'s adjoint representation in Figure 5, we found that two copies of the adjoint representation (tensor the determinant) appear. One of them has a 6-valent vertex and is thus not a largest lift. When that vertex “molts” as explained above (and the resulting hexagon is maximally inflated), one obtains the other honeycomb with the same boundary, which is a largest lift.

Proof of #2. Let m be the maximum edge-multiplicity that appears in h ; assume $m > 1$ or else we're done. Let Γ be the subgraph of $\text{supp } m_h$ of the edges with multiplicity m (and their vertices). By the assumption that β is regular, this contains none of the semi-infinite edges; it is bounded. Let x be a vertex on the boundary of the convex hull of Γ ; x is necessarily a rake or a 5-valent vertex.

Build a path γ in Γ starting at x , with first edge e , as follows. Declare the direction 60° clockwise of e to be the “forbidden” direction, and 30° counterclockwise to be “windward”.

Now traverse edges, coming to new vertices, going through crossings, turning at **Y**'s (but not into the forbidden direction), stopping when you reach another rake or 5-valent vertex. (Conceivably one might continue through a 5-valent vertex, if it is lucky enough to have two edges of multiplicity m , but we don't do this.) Once you start building this path from x, e , there are no choices, and each step carries us a positive distance in the windward direction. So the path doesn't self-intersect, and since the graph is bounded, this algorithm must terminate. By the assumption that m was the maximum edge-multiplicity, we only come into a rake or 5-valent along the edge labeled m in Lemma 10 (up to rotation and reflection).

We now attempt to simultaneously inflate all the hexagons that have completely degenerated to the vertices along γ , plus those that have collapsed to the edges connecting two vertices. Comparing this to the recipes in Lemma 10, we see that this exactly molts all the vertices, and is thus a legal combination (adding a small multiple doesn't carry us out of the cone $\text{HONEY}(\tau)$). This operation inflates some hexagons, and therefore by Lemma 9 increases the weighted perimeter, violating the largest-lift assumption as before.

We give an example of this in Figure 24, where an entire path γ molts. This illustrates how the recipes for molting at a vertex exactly fit together to give a well-defined operation on the honeycomb.

Proof of #3. Now that we know that our largest lift of a regular triple only has simple degeneracies, we can apply Lemma 2 (or rather, its strengthened version in the corollary to Theorem 1) to say that any cycle in the graph underlying the post-elision tinkertoy gives a degree of freedom. But by the assumed genericity of

w , our honeycomb should be at a *vertex* of the polytope of honeycombs lying over β , and thus have no degrees of freedom. Hence there are no cycles in the post-elision tinkertoy. \square

Honeycombs with nonsimple degeneracies can be seen in Figure 5 – but only when the bottom edge has some edges lying on top of one another, or the honeycomb is not a largest lift.

6. PROOF OF THE SATURATION CONJECTURE

We prove a general honeycomb version of the saturation conjecture. Then we derive the actual representation theory saturation conjecture from its truth for GL_n -honeycomb tinkertoys.


Theorem 3. *Let τ be a honeycomb tinkertoy such that the projection $\text{HONEY}(\tau) \rightarrow \text{BDRY}(\tau)$ is proper, and let w be a generic weighting function on the regions satisfying the properties required in section 5. Then the largest-lift map $\text{BDRY}(\tau) \rightarrow \text{HONEY}(\tau)$ is a piecewise \mathbb{Z} -linear map. Consequently, any point in $\text{BDRY}(\tau)$ assigning integer constant coordinates to the boundary edges can be extended to a lattice honeycomb.*

Proof. We already know that the largest-lift map is continuous, and linear on chambers. We will show by studying regular points in $\text{BDRY}(\tau)$ that each of these linear maps has integer coefficients. Then for any point in $\text{BDRY}(\tau)$, even nonregular, we can pick a chamber of which that point is on the boundary, and show that over that point there lies a lattice honeycomb.

If β is a regular configuration in $\text{BDRY}(\tau)$, then by Theorem 2 from section 5, the largest-lift honeycomb $l(\beta)$ has only simple degeneracies (its vertices only look like **Y**'s or crossing vertices, with edge-multiplicity 1 everywhere).

By the “elision” construction in section 2, we can regard this as a nondegenerate configuration of a simpler post-elision tinkertoy, where each crossing point is removed, and the five edges (one of length zero) replaced by the two lines going through.

In this tinkertoy, each vertex is degree 3, touching some finite and some semi-infinite edges. Consider the subgraph of finite edges, also acyclic, and inductively pull off vertices of degree 1. Each such vertex is connected to two semi-infinite edges, whose constant coordinates determine the location of the vertex. In particular the constant term on the finite edge coming out is integrally determined by those on the two semi-infinite edges – it is minus their sum.



So we can remove the two semi-infinite edges and the vertex, promoting the remaining edge to semi-infinite, and recurse. Eventually all the coordinates are integrally determined from those on the original semi-infinite edges. \square




FIGURE 25. A honeycomb tinkertoy whose largest lift shows where to find an invariant vector in a 7-fold tensor product.

(We invite the reader to see how this argument fails on the honeycomb in Figure 18.)

It may be worth noting that one can make these integral formulae very explicit. To determine the constant coordinate on a two-ended edge e somewhere in the middle of this honey-forest $l(h)$, let $e_{\mathbf{Y}}$ be the set of semi-infinite edges f such that there is a path (necessarily unique) from f to e going through e 's right-side-up \mathbf{Y} vertex. (Unless the post-elision tinkertoy is disconnected, the unique path from f to e will go either through e 's \mathbf{Y} vertex or e 's upside-down \mathbf{Y} vertex.) Then the constant coordinate on e is the sum of the constant coordinates on the outgoing boundary edges in $e_{\mathbf{Y}}$, minus the corresponding sum on the incoming edges.

Corollary (the saturation conjecture). *Let (λ, μ, ν) be a triple of dominant integral weights of $GL_n(\mathbb{C})$ such that for some $N > 0$, the tensor product $V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu}$ has a $GL_n(\mathbb{C})$ -invariant vector. Then already $V_{\lambda} \otimes V_{\mu} \otimes V_{\nu}$ has a $GL_n(\mathbb{C})$ -invariant vector.*

Proof. Let τ_n be the GL_n honeycomb tinkertoy. By Theorem 4 from Appendix 1, the fiber of the boundary-conditions map $\text{HONEY}(\tau_n) \rightarrow (\mathbb{R}^n)^3$ over the point $(N\lambda, N\mu, N\nu)$ contains a lattice point, and is therefore nonempty. Therefore the fiber over (λ, μ, ν) is also nonempty, since it is just the original fiber rescaled by $1/N$. By Proposition 2 we can apply Theorem 3, which says that this fiber contains a lattice point. Using Theorem 4 again, we find that $V_{\lambda} \otimes V_{\mu} \otimes V_{\nu}$ has a $GL_n(\mathbb{C})$ -invariant vector. \square

There is an analogous saturation conjecture for the tensor product of any number of representations, which for our purposes we can state as follows. Let $\{\lambda_i\}_{i=1, \dots, m}$ be a collection of dominant weights such that for some large N , the tensor product $\otimes_i V_{N\lambda_i}$ has a $GL_n(\mathbb{C})$ -invariant vector. Then the same is true when N is replaced by 1.

An earlier version of this paper had a technically unpleasant proof of the general saturation conjecture. We omit the details, because since writing this paper, Andrei Zelevinsky has shown us how to derive the general saturation result from the case already proven, via standard arguments with the Littlewood-Richardson rule [Ze]. The basic idea of the omitted proof is indicated in Figure 25, which shows a honeycomb tinkertoy whose largest lifts give witnesses to saturation in the case of a seven-fold tensor product.

As long as the inputs λ_i are pulled apart sufficiently ($x \rightarrow \infty$), a configuration of this big honeycomb tinkertoy corresponds 1:1 to a ‘coherent’ set of $(0, n, 0, n, 0, n)$ honeycombs (they are far enough apart to necessarily not overlap when glued together into a configuration of the big honeycomb tinkertoy). That and repeated application of the isomorphism

$$(A \otimes B)^{GL_n(\mathbb{C})} \cong \sum_{\mu} (A \otimes V_{\mu})^{GL_n(\mathbb{C})} \otimes (V_{\mu}^* \otimes B)^{GL_n(\mathbb{C})}$$

(where μ runs over all irreducible representations of $GL_n(\mathbb{C})$) let us locate a lattice point in the analogous Berenstein-Zelevinsky polytope.

7. A SATURATION CONJECTURE FOR OTHER GROUPS

One can phrase a naïve saturation conjecture for other groups G : for any triple (λ, μ, ν) of dominant weights for G , if there exists a number N such that

$$(V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu})^G > 0, \quad \text{then} \quad (V_{\lambda} \otimes V_{\mu} \otimes V_{\nu})^G > 0.$$

However, this conjecture is false, with counterexamples reported in [E]. A clue is provided by the fact that it is already false for $G = SL_n(\mathbb{C})$!

Of course, the $SL_n(\mathbb{C})$ case is not so different from the $GL_n(\mathbb{C})$ case, where saturation holds, and so is easily patched up; we must ask also that the representation $\lambda + \mu + \nu$ of the torus annihilate the center of $SL_n(\mathbb{C})$. (This is no longer a linear condition on $\lambda + \mu + \nu$, as it was for $GL_n(\mathbb{C})$, exactly because $SL_n(\mathbb{C})$ ’s center is not connected.)

Conjecture. *Let G be a connected complex semisimple Lie group with maximal torus T , (λ, μ, ν) a triple of dominant weights for G , and N a positive number such that*

$$(V_{N\lambda} \otimes V_{N\mu} \otimes V_{N\nu})^G > 0.$$

Then if $\lambda + \mu + \nu$ annihilates all elements of T with semisimple centralizer,

$$(V_{\lambda} \otimes V_{\mu} \otimes V_{\nu})^G > 0.$$

We now explain the geometric motivation for this conjecture. Since our only goal is to make the conjecture plausible we do not waste space on full proofs of the statements made here.

The space $(V_{\lambda} \otimes V_{\mu} \otimes V_{\nu})^G$ can be thought of as the space of sections of a sheaf on the GIT quotient $(G/B)^3 // G$ (as explained in the introduction, in the case of $G = GL_n(\mathbb{C})$). However, this sheaf may not necessarily be a line bundle; it can have singularities that get worse at orbifold points of the quotient, and any section is required to vanish at the singularities. The condition that $\lambda + \mu + \nu$ annihilate the center exactly guarantees that this sheaf be a line bundle generically, so is certainly necessary for the existence of nonvanishing sections.

However, if the sheaf has singularities along which any section must vanish, it stands to reason that global sections are less likely to exist. The condition in the conjecture is exactly equivalent to asking that the sheaf be a line bundle globally, therefore to have a better chance to have sections.

To be sure, this conjecture is not nearly as satisfactory as the result for $GL_n(\mathbb{C})$ (primarily because it's not necessary, only claimed to be sufficient), but the situation for other groups seems to be inherently less clean.

1. APPENDIX: THE EQUIVALENCE OF GL_n -HONEYCOMBS WITH A DEFINITION OF BERENSTEIN-ZELEVINSKY PATTERNS

Let η be a hexagon with 120° angles and two vertical edges. Define the **torsion** of η to be the length of the left edge minus that of the right edge.

Proposition 3. *The torsion of η and that of each 120° rotation of η agree.*

Proof. For a regular hexagon they are all zero. If one translates one edge of η out from the center, the edge shrinks and its two neighboring edges grow, keeping the torsions equal. Any position of the hexagon can be achieved by composing such translations. □

Let h be a GL_n -honeycomb. We will assign a number to each region in the GL_n honeycomb tinkertoy τ_n , other than the sectors at the three corners, using h . This will turn out to be a Berenstein-Zelevinsky pattern.

1. Each hexagon is assigned its torsion.
2. Each semi-infinite wedge on the NE long edge of the honeycomb is assigned the length of its west edge.
3. Each semi-infinite wedge on the NW long edge of the honeycomb is assigned the length of its SE edge.
4. Each semi-infinite wedge on the bottom long edge of the honeycomb is assigned the length of its NE edge.

(This is set up so as to be 120° -rotation invariant.)

For any region not on the NW long edge, the sum of the region-entries at and to the right of that point telescopes to the length of an edge, necessarily nonnegative. (Likewise for 120° rotations.)

To determine the sum across an entire row is a little trickier. We need to relate the length of an edge to the constant coordinates on neighboring edges. Rotate the edge to align it with the finite edge in Figure 26.

So the sum across an entire left-right row, with a semi-infinite wedge at the left end, is

1. the sum of the lengths of the two finite edges of that wedge;
2. the difference of the constant coordinates of the semi-infinite edges of that wedge.

(By 120° -rotational symmetry the same is true for sums in other directions.)




FIGURE 26. Formula for the length of an edge. The edges in the honeycomb are labeled with their constant coordinates.

In particular, if the constant coordinates of the semi-infinite edges are interpreted as the coefficients λ_i, μ_i, ν_i of three dominant weights (in nonincreasing order) as in the rest of the paper, the labeling of the regions exactly matches the definition of **Berenstein-Zelevinsky pattern** as given in [Ze] (only one of many realizations, others to be found in the original [BZ]).

The central theorem in [BZ] gives a formula for Littlewood-Richardson coefficients as the number of BZ patterns with given boundary values. Their formulation is more suited to SL_n than GL_n calculations, and as such, they need to include a caveat that BZ patterns count the LR coefficient only if the sum of the three weights is in the root lattice of SL_n . (If the sum is not in the root lattice, the LR coefficient is obviously zero, but there are still likely to be many BZ patterns which now have no known representation-theoretic meaning.) In the GL_n -adapted formulation of this paper this caveat does not appear.

For us, the BZ theorem reads as follows:

Theorem 4. *Let λ, μ, ν be a triple of dominant weights of $GL_n(\mathbb{C})$. Then the number of lattice GL_n -honeycombs whose semi-infinite edges, indexed clockwise from the southwest, have constant coordinates $\lambda_1, \dots, \lambda_n, \mu_1, \dots, \mu_n, \nu_1, \dots, \nu_n$ is the corresponding Littlewood-Richardson coefficient $\dim(V_\lambda \otimes V_\mu \otimes V_\nu)^{GL_n(\mathbb{C})}$.*

2. APPENDIX: THE HIVE MODEL

In this section we introduce another model of the points in the Berenstein-Zelevinsky cone, much closer to the BZ models in feel, that has the honeycomb-like property of having only “local” inequalities. It is not strictly necessary for the logic of the paper, but is very useful as an alternate model; this is particularly true if one wants to actually count tensor product multiplicities, rather than merely prove them positive. The paper [Bu] exposing our work takes this model as the fundamental one.

As was mentioned elsewhere, the primary advantage of the honeycomb model over the original BZ models is the naturality of the “overlay” operation; in this paper this is only used in a sort of local way, when we elide simple degeneracies. But this comes with a cost – the linear structure on the space of honeycombs is a bit difficult to see geometrically. In addition, the degrees of freedom of the honeycomb are distinctly more obscure than in the BZ models.¹⁸ (And then there is the typesetting problem.)

Given τ a honeycomb tinkertoy, recall the dual graph $D(\tau)$ defined in section 3. Define the vector space $\mathbf{hive}(\tau)$ to be labelings of the vertices of $D(\tau)$ by real numbers. This vector space naturally contains the lattice of integer labelings.

Recall that the embedded graph $D(\tau)$ is a collection of triangles, which we refer to as the **hive triangles**. Of most interest to us are the rhombi formed by pairs of adjacent hive triangles. Each such rhombus has two **acute vertices** and two **obtuse vertices**. There are three possible directions a rhombus may face (see Figure 27).

Each rhombus $\rho \subseteq H$ gives a functional on $\mathbf{hive}(\tau)$, defined as the sum at the obtuse vertices minus the sum at the acute vertices (as seen in Figure 28). This gives a **rhombus inequality**, asking that the rhombus functional be nonnegative.

¹⁸In some sense, though, the strength of the honeycomb model as used in this paper is that one can study the degrees of freedom left *while keeping some of the inequalities pressed*, which is not so easy to do in the BZ models.




FIGURE 27. The dual graph $D(\tau)$ corresponding to the GL_6 honeycomb tinkertoy, a little hive triangle, and a rhombus in each orientation.




FIGURE 28. The dot-product picture of a rhombus inequality.

The cone in $\mathbf{hive}(\tau)$ satisfying all these inequalities we denote $\mathbf{HIVE}(\tau)$, and we will call its elements τ -hives.

Proposition 4. *Let τ be a honeycomb tinkertoy. Then there is a \mathbb{Z} -linear correspondence between configurations of τ , and τ -hives whose leftmost top entry is zero, in such a way that the constant coordinates of the boundary edges of the honeycomb are differences of boundary entries on the corresponding τ -hive.*

Proof. Start with a configuration h of τ . We label the vertices in $D(\tau)$ inductively, starting with a zero in the leftmost top entry, and filling in as follows: whenever we move southwest or east, we increase the value by the constant coordinate of the edge crossed; southeast, we decrease by that constant coordinate.

Our first worry is that different paths will cause us to try to fill different numbers in the same hexagon. That this doesn't happen is a simple consequence of the sum-equals-zero property at a vertex of the honeycomb.

Second, we need to know that the result is a hive. Not surprisingly, the rhombus inequalities are equivalent to the edge lengths being nonnegative.

Lastly, since we define the hive entries by inductively adding up constant coordinates of edges, the boundary of the hive naturally ends up being the partial sums of those constant coordinates. (The sum-equals-zero property is involved in seeing this for some of the boundary edges.) \square

Combining this with the theorem in the appendix relating GL_n -honeycombs to Littlewood-Richardson coefficients, we find that if λ, μ, ν are integral, the number of hives with boundary formed from partial sums of λ, μ, ν is a Littlewood-Richardson coefficient. (In [Bu], there is given a simple bijection between hives and a standard formulation of the Littlewood-Richardson rule.)

There is a pleasant geometric way to interpret the rhombus inequalities. Extend the hive to a piecewise linear function, affine-linear on each little hive triangle. Then the rhombus inequalities state that this function is convex. Each rhombus *equality* says that the function is actually linear across the boundary down the middle of the rhombus, i.e. that the regions on which the function is affine-linear are larger than just the little hive triangles.

In this way, the set of tight rhombus inequalities determines a certain decomposition of the convex region in $\mathbb{R}_{\sum=0}^3$ bearing $D(\tau)$ into regions (which we dub “flatspaces” due to the geometric interpretation above). This is exactly the decomposition into the regions of the degeneracy graph from section 3. We mentioned there that the degeneracy graph remembers only the “combinatorial information” about a honeycomb; we see now that it is the hive that finishes the job.

Again, we refer readers to the honeycomb/hive applet to see these hives and convex graphs in action, at

<http://www.alumni.caltech.edu/~allenk/java/honeycombs.html>.

REFERENCES

- [Bu] A. Buch, The saturation conjecture (after A. Knutson and T. Tao), notes from a talk at Berkeley, September 1998.
- [BS] L. Billera, B. Sturmfels, Fiber polytopes, *Annals of Math.*, **135** (1992), no. 3, 527–549. MR **93e**:52019
- [BZ] A. Berenstein, A. Zelevinsky, Triple multiplicities for $sl(r+1)$ and the spectrum of the exterior algebra of the adjoint representation, *J. Alg. Comb.*, **1** (1992), 7–22. MR **93h**:17012
- [E] A.G. Elashvili, Invariant algebras, *Advances in Soviet Math.*, **8** (1992), 57–64. MR **93c**:17013
- [F] W. Fulton, Eigenvalues of sums of Hermitian matrices (after A. Klyachko), Séminaire Bourbaki (1998).
- [FH] W. Fulton, J. Harris, Representation theory, Springer-Verlag (1991). MR **93a**:20069
- [GP] O. Gleizer, A. Postnikov, Littlewood-Richardson coefficients via Yang-Baxter equation, in preparation.
- [GZ] V. Guillemin, C. Zara, Equivariant de Rham theory and graphs, [math.DG/9808135](https://arxiv.org/abs/math/9808135).
- [H] A. Horn, Eigenvalues of sums of Hermitian matrices, *Pacific J. Math.*, **12** (1962), 225–241. MR **25**:3941
- [Hon2] A. Knutson, T. Tao, C. Woodward, The honeycomb model of $GL_n(\mathbb{C})$ tensor products II: Facets of the L-R cone, in preparation.
- [KSZ] M. Kapranov, B. Sturmfels, A. Zelevinsky, Quotients of toric varieties. *Math. Annalen*, **290** (1991), no. 4, 643–655. MR **92g**:14050
- [KI] A.A. Klyachko, Stable vector bundles and Hermitian operators, *IGM, University of Marne-la-Vallee preprint* (1994).
- [KMP] S. Kumar, Proof of the Parthasarathy-Ranga Rao-Varadarajan conjecture, *Invent. math.*, **93** (1988), no. 1, 117–130. MR **89j**:17009
- [M] O. Mathieu, Construction d’un groupe de Kac-Moody et applications. *Compositio Math.* **69** (1989), no. 1, 37–60. MR **90f**:17012
- [P] P. Polo, Variétés de Schubert et filtrations excellentes. *Astérisque*. **10-11** (1989) 281–311. MR **91b**:20056
- [Ku] G. Kuperberg, Spiders for rank two Lie algebras, [math.QA/9712143](https://arxiv.org/abs/math/9712143), *Comm. Math. Phys.*, **180** (1996), no. 1, 109–151. MR **97f**:17005
- [MFK] D. Mumford, J. Fogarty, F. Kirwan, Geometric invariant theory. Third edition. Chapter 8. *Ergebnisse der Mathematik und ihrer Grenzgebiete*, Springer-Verlag, 1994. MR **95m**:14012

- [Ze] A. Zelevinsky, Littlewood-Richardson semigroups, [math.CO/9704228](#).
[Zi] G. Ziegler, Lectures on polytopes, *Graduate Texts in Mathematics*, 152. Springer-Verlag, 1995. MR [96a:52011](#)

DEPARTMENT OF MATHEMATICS, BRANDEIS UNIVERSITY, WALTHAM, MASSACHUSETTS 02254
Current address: Department of Mathematics, University of California Berkeley, Berkeley, California 94720-3840
E-mail address: allenk@alumni.caltech.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA LOS ANGELES, LOS ANGELES, CALIFORNIA 90095-1555
E-mail address: tao@math.ucla.edu