HIVES
NOTES FOR MATH 261, FALL 2001

ALLEN KNUTSON

Consider the triangle $x + y + z = n, x, y, z \geq 0$. This has $\binom{n+2}{2}$ integer points; call this set $\mathtt{tri}_n$. We will draw it in the plane with $(n, 0, 0)$ at the lower left, $(0, n, 0)$ at the top, and $(0, 0, n)$ in the lower right.

A **hive order** n, or n**-hive**, is a function $h : \mathtt{tri}_n \to \mathbb{Z}$ satisfying certain inequalities. Here are four equivalent ways to state them:

1. On each unit rhombus in the triangle, the sum across the short diagonal is greater than or equal to the sum across the long diagonal.
2. $$h(x+1, y, z+1) + h(x, y+1, z+1) \geq h(x+1, y+1, z) + h(x, y, z+2)$$
   when these four points are all in $\mathtt{tri}_n$, and likewise for the 120° and 240° rotations of the hive.
3. If you extend h to a real-valued function on the solid triangle by making it linear on each little triangle $\overline{(x+1, y, z)\ (x, y+1, z)\ (x, y, z+1)}$, h is convex.
4. If you have two $0, \pm 1$-vectors $\vec{v}, \vec{w}$ of sum zero, with dot product 1, then
   $$\mathtt{diff}_{\vec{v}} \cdot \mathtt{diff}_{\vec{w}} h \geq 0$$
   wherever it is defined.

(Note that the definition also makes sense for real-valued functions, in which case we will speak of a **real hive**.)

Call these inequalities the **rhombus inequalities** on a hive. They naturally come in three families, according to the orientation of the rhombus.

Proposition. *Let $a_0, a_1, \ldots, a_n$ be the numbers on one side of a hive. Then a is convex, i.e. $a_i \geq \frac{1}{2}(a_{i-1} + a_{i+1})$. Put another way, the list $(a_i - a_{i+1})$ is a dominant weight for $GL_n(\mathbb{C})$.*

Proof. There are two rhombi with an obtuse vertex at a_i. Adding the two corresponding rhombus inequalities, we get the desired result. □

We will mostly be interested in hives whose lower left entry is zero. Let $\mathtt{HIVE}_{\lambda\mu\nu}$ denote the set of hives (with lower left entry zero) such that the differences in the boundary entries, going around counterclockwise, gives the dominant weights λ, μ, ν for $GL_n(\mathbb{C})$. Since we get back to zero when we're done, this requires $\sum_i(\lambda_i + \mu_i + \nu_i) = 0$ for $\mathtt{HIVE}_{\lambda\mu\nu}$ to be nonempty.

A side note (we could have made at any time). Given a dominant weight λ, let λ^* be $-\lambda$ written in reverse order (so it is again weakly decreasing). Then $V_{\lambda^*} \cong (V_\lambda)^*$. Proof: the highest weight of V_λ is minus the lowest weight of $(V_\lambda)^*$. Knowing the lowest weight, we can figure out the highest weight by hitting it with the permutation that reverses everything.

There is a close relation of hives to G-C patterns, which we motivate with the following result:

Proposition. *Let λ, π, ν be dominant weights of $GL_n(\mathbb{C})$ such that $\forall i, \pi_i - \pi_{i+1} \geq \lambda_1 - \lambda_n$. Then the multiplicity of V_π in $V_\lambda \otimes V_\nu$ reduces to a weight multiplicity, of the $\pi - \nu$ weight in V_λ.*

Proof. Ordinarily the Steinberg tensor product formula would only give this as an alternating sum, $\sum_{S_n} (-1)^w$ times the $(\pi - w(\nu + \rho) + \rho)$-weight multiplicity in V_λ. In this extreme case, however, one can check that all these terms vanish except for $w = 1$. □

Our goal in this section is to prove that $\mathtt{HIVE}_{\lambda\mu\nu}$ computes the dimension of the invariant space of $V_\lambda \otimes V_\mu \otimes V_\nu$. For the moment, we content ourselves with this.

Theorem. *Give a hive $h \in \mathtt{HIVE}_{\lambda\mu\nu}$, let $\partial h : \mathtt{tri}_{n-1} \to \mathbb{Z}$ be defined by $\partial h(x, y, z) = h(x, y, z + 1) - h(x + 1, y, z)$. Graphically, this subtracts from each entry on the hive the entry to the left (unless there is no left; the upper left side is thrown out).*

Then ∂ gives an injection from $\mathtt{HIVE}_{\lambda\mu\nu}$ to the set of G-C patterns with λ on the bottom, and having weight $\mu^* - \nu$.

Proof. Obviously we get something that looks like a G-C pattern, i.e. an array of $\binom{n+1}{2}$ numbers in a triangle. By definition, it has λ across the bottom.

To be a G-C pattern, it has to satisfy the interspersing inequalities. These follow from the two families of rhombus inequalities having horizontal edges.

Next, we consider the weight. Since the entries in ∂h are constructed as differences in h, if we sum a row of ∂h, we get the rightmost entry in h minus the leftmost.

Let C be the top entry in h (in fact it is $\sum_i \lambda_i + \mu_i = -\sum_i \nu_i$). Then the leftmost entry in the $(i + 1)$st row of h is $C + \nu_1 + \nu_2 + \ldots + \nu_i$, whereas the rightmost is $C - \mu_n - \mu_{n-1} - \ldots - \mu_{n-i+1} = C + \mu_1^* + \mu_2^* + \ldots + \mu_i^*$.

Therefore the sum of the ith row of ∂h is $(\mu_1^* + \mu_2^* + \ldots + \mu_i^*) - (\nu_1 + \nu_2 + \ldots + \nu_i)$. So the weight of the Gel′fand-Cetlin pattern ∂h, which is defined as the *differences* of the row sums, is $\mu^* - \nu$.

So far we know that ∂ lands in the advertised target. Now why is it injective? Basically, we know the "initial conditions" of the hive (the labeling on the left side), and from that and the differences ∂h we can reconstruct h by partial-summing ∂h. □

What keeps it from being bijective – didn't we just construct an inverse? Not quite, because the partially-summed guy may not satisfy the third family of rhombus inequalities. One can show, though, that if each $\mu_i - \mu_{i+1} \geq \lambda_1 - \lambda_n$, this map ∂ is a (linear) bijection.

We had a homework problem (the one about half-Gel′fand-Cetlin patterns) that showed that the Kostant partition function is a certain limiting case of weight multiplicities, essentially by saying that satisfying two families of rhombus inequalities can degenerate to satisfying one, if the other becomes automatic.

The theorem above shows that weight multiplicities are themselves a limiting case of hives, in the case that satisfying three families of rhombus inequalities degenerates to satisfying two.

Let $\mathrm{Rep}(GL_n(\mathbb{C}))$ denote the ring of (formal differences of algebraic finite-dimensional) representations of $GL_n(\mathbb{C})$, with addition and multiplication coming from direct sum and tensor product of representations. Then $\mathrm{Rep}(GL_n(\mathbb{C}))$ has a canonical basis $\{[V_\lambda]\}$, the irreps, indexed by the set of dominant weights. (The "$[]$" are only there to maintain a proper distinction between an actual representation V_λ and the corresponding element of $\mathrm{Rep}(GL_n(\mathbb{C}))$, which is really an isomorphism class.)

Let $Q_m \subset \mathrm{Rep}(GL_n(\mathbb{C}))$, $m \in \mathbb{Z}$, be spanned by the $[V_\lambda]$ with $\sum_i \lambda_i = m$.

Lemma. $\mathrm{Rep}(GL_n(\mathbb{C})) = \oplus_{m\in\mathbb{Z}} Q_m$, *and* $Q_m Q'_m \leq Q_{m+m'}$.

Proof. The first statement is obvious – we've just partitioned the basis. For the second, recall that $\sum_i \lambda_i$ is just the weight of the action of the center of $GL_n(\mathbb{C})$ (the scalar matrices) on V_λ. So on $V_\lambda \otimes V_\mu$, the center acts with weight $\sum_i \lambda_i + \mu_i$. □

This gave a gradation of $\mathrm{Rep}(GL_n(\mathbb{C}))$. We'll need also a somewhat trickier filtration. Define a partial order $\leq$ on weights, by $\mu \leq \lambda$ if $\mu - \lambda$ is a sum of negative roots. Equivalently, $\sum_i \mu_i = \sum_i \lambda_i$, but each partial sum $\sum_{i=1}^k \mu_i$ is bounded above by the corresponding partial sum of λ_i.

For each dominant weight λ, let R_λ be the subspace of $\mathrm{Rep}(GL_n(\mathbb{C}))$ spanned by the set of $[V_\mu]$, for $\mu \leq \lambda$.

Lemma. $R_\lambda R_\mu \subseteq R_{\lambda+\mu}$.

Proof. This follows from a very lame version of the Steinberg multiplicity formula, lame enough that we construct it from scratch here.

Let $[V_{\lambda'}], [V_{\mu'}]$ be two basis elements of R_λ, R_μ (i.e. irreps with $\lambda' \leq \lambda$, $\mu' \leq \mu$). Then $V_{\lambda'}$ is strongly dominated by λ'. In particular, every weight in $V_{\lambda'}$ is $\leq \lambda'$. Similarly μ'.

From there, we get that every weight of $V_{\lambda'} \otimes V_{\mu'}$ is $\leq \lambda + \mu$.

Now when we break up that weight diagram by stripping off irreducibles, we only use irreps whose high weights are $\leq \lambda + \mu$. These are the basis vectors in $R_{\lambda+\mu}$. □

Theorem. *The ring* $\mathrm{Rep}(GL_n(\mathbb{C})) \cong \mathbb{Z}[a_1, a_2, \ldots, a_n, a_n^{-1}]$, *where* $a_k = [\mathrm{Alt}^k(\mathbb{C}^n)]$.

Proof. Certainly we have a map ϕ from the second to the first, taking a_k to $[\mathrm{Alt}^k(\mathbb{C}^n)]$. To get a representation strongly dominated by λ, use $\phi(a_1^{\lambda_1-\lambda_2} a_2^{\lambda_2-\lambda_3} \ldots a_{n-1}^{\lambda_{n-1}-\lambda_n} a_n^{\lambda_n})$. We already determined a while ago that any list of representations X_λ indexed and strongly dominated by dominant weights λ gives a basis. So this shows ϕ is onto.

Note also that $\phi(a_1^{\lambda_1-\lambda_2} a_2^{\lambda_2-\lambda_3} \ldots a_{n-1}^{\lambda_{n-1}-\lambda_n} a_n^{\lambda_n}) \in Q_{\sum_i \lambda_i}$.

If it's not injective, then there's some polynomial $P(a_1, \ldots, a_n, a_n^{-1})$ such that $\phi(P) = 0$. Writing $P = \sum_m P_m$ by grouping the different monomials in P by the Q_m they go into, we see that each $\phi(P_m)$ must be zero separately. So we may as well assume that $P = P_m$ for some fixed m.

Let $c \prod_i a_i^{p_i}$ be a monomial in P with largest p_1, then largest p_2 (among those with largest p_1), and so on. Define λ by $\lambda_i = \sum_{j \geq i} m_j$. It follows that for each monomial m in P, $\phi(m) \in R_\lambda$.

But in fact, for each monomial m in P other than $c\prod_i a_i^{p_i}$, $\phi(m) \in R_{\lambda'}$ for some λ' strictly $< \lambda$. Therefore the image of $\phi(P)$ in $R_\lambda / \sum_{\lambda'<\lambda} R_{\lambda'}$ is nonzero, contradiction. □

So this ring is rather boring as just a ring. What's interesting about it is the canonical basis (of irreps), and the product structure on that basis (rewriting a product as the sum of basis elements).

2. The hive ring (assuming it's associative)

Define a (nonassociative?) ring H_n with a basis $\{b_\lambda\}$ indexed by $GL_n(\mathbb{C})$'s dominant weights, and the multiplicative structure

$$b_\lambda b_\mu = \sum_\nu \#\mathtt{HIVE}_{\lambda\mu\nu^*}\, b_\nu.$$

The right way to think about this set of hives is that the differences on the NW side are λ, on the NE side are μ, and on the S side are ν, *all read left to right*.

The most technical thing we will have to prove is

Theorem. *The hive ring* H_n *is actually associative.*

We'll prove this in the next section. The goal now is to show that the correspondence $b_\lambda \mapsto [V_\lambda]$ gives a ring isomorphism $H_n \cong \mathrm{Rep}(GL_n(\mathbb{C}))$, and therefore that counting hives computes tensor product multiplicities.

Define R_λ as we did for $\mathrm{Rep}(GL_n(\mathbb{C}))$. We prove that it gives a filtration of this ring too:

Lemma. $R_\lambda R_\mu \subseteq R_{\lambda+\mu}$. *Also,* $b_\lambda b_\mu$ *contains* $b_{\lambda+\mu}$ *with coefficient* 1.

Proof. Consider hives with labels (the partial sums of) λ on the NW, μ on the NE, and ν on the S, all read left-to-right. Then $\sum_i \nu_i = \sum_i \lambda_i + \sum_i \mu_i$ automatically. We want to show that $\nu \leq \lambda + \mu$ in the partial order on weights, and that equality is only the easy part of it.

We first mention a sublemma: given any lattice parallelogram in a hive, with the edges oriented in coordinate directions (of which there are three, not two), the sum of the terms at the obtuse vertices beats the sum of the terms at the acute vertices. Proof: the parallelogram breaks up into a sum of rhombi. Adding up the rhombus inequalities, all dependence on the internal numbers cancels and only the corners remain.

Now consider such parallelograms with one acute vertex at the very top, and the other at the ith point along the bottom. then this parallelogram inequality says that

$$\sum_{k=1}^{i} \nu_i \leq \sum_{k=1}^{i} \lambda_i + \sum_{k=1}^{i} \mu_i.$$

This is the hard half of establishing $\nu \leq \lambda + \mu$.

So the only terms in $b_\lambda b_\mu$ are b_ν with $\nu \leq \lambda + \mu$. This shows the containment presented.

For the second statement, make a hive by setting each entry equal to the sum of the obtuse vertices, minus the top vertex $\sum \lambda_i$, of the parallelogram connecting that entry to the top entry. This is easily checked to be a hive.

To see that it's the only one with bottom $\lambda + \mu$, write the parallelogram inequalities that bounded the bottom entries as the sum of two smaller parallelogram inequalities. For the whole inequality to be pressed, the individual ones must be pressed. □

Theorem. *The hive ring* $H_n \cong \mathbb{Z}[a_1, \dots, a_n, a_n^{-1}]$, *where* $a_i = b_{(1,\dots,1,0,\dots,0)}$, $a_n^{-1} = b_{(-1,-1,\dots,-1)}$.

Proof. This is essentially the same proof as in the $Rep(GL_n(\mathbb{C}))$ case. If some polynomial dies, look at its leading term, and show that that doesn't die in some $R_\lambda / \sum_{\lambda'<\lambda} R_{\lambda'}$, contradiction. □

The **Pieri rule** says that

$$V_\lambda \otimes \mathrm{Alt}^k(\mathbb{C}^n) \cong \sum_{\substack{\mu \text{ a weight of } \mathrm{Alt}^k \\ \lambda+\mu \text{ dominant}}} V_{\lambda+\mu}.$$

It is easy to prove from the Steinberg tensor product rule (homework problem).

Lemma. *The product* $b_\lambda b_{1,\dots,1,0,\dots,0}$ *in the hive ring also decomposes á la Pieri.*

Proof. Also a homework problem. □

From these facts and associativity, we get the full result:

Theorem. *The linear isomorphism* $\phi : H_n \cong Rep(GL_n(\mathbb{C}))$, *taking each* b_λ *to* $[V_\lambda]$, *is a ring isomorphism.*

Proof. We want to show that $\phi(yx) = \phi(y)\phi(x)$. By linearity, it's enough to show it for y a basis element b_λ.

The Pieri rule being true in both rings then tells us that this equation does hold if x is a generator, $b_{(1,\dots,1,0,\dots,0)}$ or $b_{(-1,-1,\dots,-1)}$.

More generally, let $x = b_{\mu_1} b_{\mu_2} \dots b_{\mu_l}$ be a product of l generators. Then

$$\phi(b_\lambda(b_{\mu_1} b_{\mu_2} \dots b_{\mu_l})) = \phi((b_\lambda b_{\mu_1} b_{\mu_2} \dots b_{\mu_{l-1}}) b_{\mu_l}) = \phi(b_\lambda b_{\mu_1} b_{\mu_2} \dots b_{\mu_{l-1}})\phi(b_{\mu_l})$$

and induction on l takes care of the rest.

Then the statement that these *are* generators is exactly that H_n is linearly spanned by monomials in the generators. So $\phi(b_\lambda x) = \phi(b_\lambda)\phi(x)$ for any λ and x, and we're done. □

3. The hive ring is indeed associative

First off, what's the equation we're trying to prove? Let $h^\sigma_{\lambda\mu}$ be the number of hives with λ, μ across the top two sides, σ across the bottom, all read left-to-right. Then

$$(b_\lambda b_\mu) b_\nu = \sum_\sigma h^\sigma_{\lambda\mu} b_\sigma b_\nu = \sum_\sigma \sum_\pi h^\sigma_{\lambda\mu} h^\pi_{\sigma\nu} b_\pi$$

whereas

$$b_\lambda (b_\mu b_\nu) = \sum_\tau b_\lambda h^\tau_{\mu\nu} b_\tau = \sum_\tau \sum_\pi h^\tau_{\mu\nu} h^\pi_{\lambda\tau} b_\pi$$

Extracting coefficents of b_π, we get

$$\sum_\sigma h^\sigma_{\lambda\mu} h^\pi_{\sigma\nu} = \sum_\tau h^\tau_{\mu\nu} h^\pi_{\lambda\tau}$$

Consider a tetrahedron balanced perfectly on an edge, from above; the boundary of what you see is a square. Label the edges of this square (starting from one of the lower two vertices and going clockwise) with the partial sums of λ, μ, ν, π^*. If the bottom edge is labeled σ, then the number of ways of labeling the lower two faces with hives is $h^{\sigma}_{\lambda\mu} h^{\pi}_{\sigma\nu}$. Without fixing the labeling, it's $\sum_\sigma h^{\sigma}_{\lambda\mu} h^{\pi}_{\sigma\nu}$. The corresponding statement for the top two faces gives the other sum.

Proposition. *There is a (piecewise linear) bijection between ways of labeling the upper two faces of this tetrahedron with a pair of hives and ways of labeling the lower two faces, with given fixed labels λ, μ, ν, π^* around the four non-horizontal edges.*

This proof was found by Chris Woodward, in the context of honeycombs.

Proof. This tetrahedron of size n breaks up into little tetrahedra, little upside-down tetrahedra, and octahedra (think about the $n = 2$ case). We will excavate it from the top, helicoptering out pieces only when everything above them is already out of the way.

Whenever we remove a little tetrahedron, we don't expose any new lattice points. Whenever we remove an octahedron, though, one of the old vertices (a local max) goes with it and a new one becomes visible (a local min). As we go, we label the vertices exposed according to the following formula:

$$e' := \max(a + c, b + d) - e$$

where e was the label at the top, and a, b, c, d the labels around the equatorial square.

When we're done, we have labeled the bottom two faces. The process obviously provides its own inverse.

It remains to see that what we get on the bottom is a pair of hives, i.e. satisfies the rhombus inequalities. We claim that *every* unit rhombus in the tetrahedron gives a true rhombus inequality.

Say we've partially excavated, and every rhombus above the level so far dug has satisfied this inequality. Now we extract a piece; this exposes some new rhombi that we need to check.

The $n = 2$ case. We remove the top two tetrahedra, then the octahedron, then a bottom tetrahedron. From the top, we see the labels

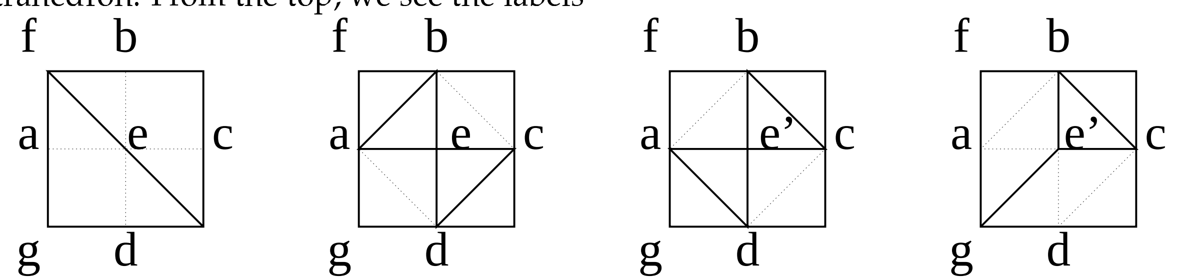

The second move exposes the rhombus with obtuse vertices a, b, acute $f, e' = \max(a + c, b + d) - e$. We want to show that

$$a + b \geq f + \max(a + c, b + d) - e$$

or equivalently

$$a + b \geq f + a + c - e, a + b \geq f + b + d - e$$

which follow from the $b + e \geq f + c, a + e \geq d + f$ inequalities on the top.

The third move exposes the rhombus with obtuse vertices a, e', acute b, g. We want to show that

$$a + \max(a + c, b + d) - e \geq b + g$$

so it's enough to show one of them: $a+b+d-e \geq b+g$. This follows from $a+d \geq e+g$ on the top.

The general case. Any rhombus exposed fits into a size 2 tetrahedron, so we just have to apply the $n=2$ case over and over. ☐

Here's an example, computing the tensor square of the $(2,1,0)$ representation of $GL_3(\mathbb{C})$.

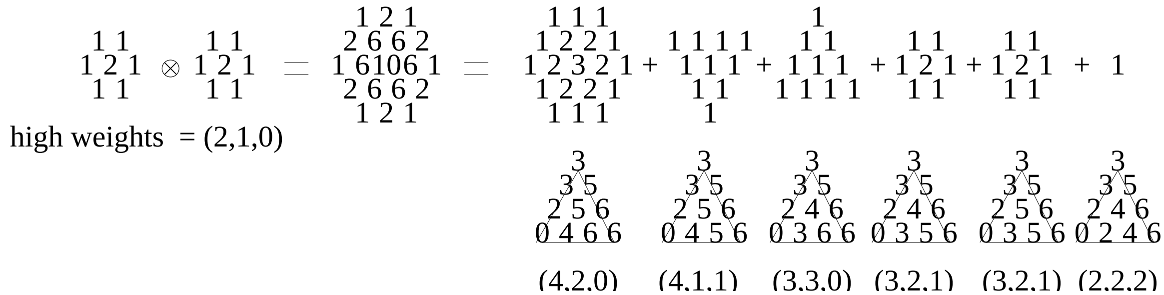

Note that this representation restricts to the adjoint representation of $SL_3(\mathbb{C})$. The bracket on that Lie algebra $\mathfrak{sl}_3$ gives an equivariant map from $\mathrm{Alt}^2\mathfrak{sl}_3$ to $\mathfrak{sl}_3$, which is why we're not surprised to find a copy of the adjoint rep inside its tensor square. (In fact we find two.) The trivial rep is in there because of the invariant form $(X, Y) \mapsto \mathrm{Tr}\,(XY)$ on $\mathfrak{sl}_3$.