Honeycombs and Sums of Hermitian Matrices

Allen Knutson and Terence Tao

In 1912 Hermann Weyl [W] posed the following problem: given the eigenvalues of two $n \times n$ Hermitian matrices A and B, how does one determine all the possible sets of eigenvalues of the sum $A+B$? When $n=1$, the eigenvalue of $A+B$ is of course just the sum of the eigenvalue of A and the eigenvalue of B, but the answer is more complicated in higher dimensions. Weyl's partial answers to this problem have since had many direct applications to perturbation theory, quantum measurement theory, and the spectral theory of self-adjoint operators. The purpose of this article is to describe the complete resolution to this problem, based on recent breakthroughs [Kl], [HR], [KT], [KTW].

To standardize the notation, we shall always write the eigenvalues of an $n \times n$ Hermitian matrix as a weakly decreasing n-tuple $\lambda = (\lambda_1 \geq \ldots \geq \lambda_n)$ of real numbers. Thus, for instance, the eigenvalues of diag$(3,2,5,3)$ are $(5,3,3,2)$.

To illustrate Weyl's problem, suppose that $n=2$ and that A, B have eigenvalues $(3,0)$ and $(5,0)$ respectively. Then one can easily verify that $A+B$ can have eigenvalues $(8,0)$ or $(5,3)$ or, more generally, $(8-a, a)$ for any $0 \leq a \leq 3$. This turns out to be the complete set of possibilities; $A+B$ cannot have eigenvalues $(9,-1)$ or $(7,0)$ or $(4,4)$, etc.

Allen Knutson is professor of mathematics at the University of California Berkeley. His e-mail address is allenk@math.berkeley.edu. *His research was partially conducted for the Clay Mathematics Institute.*

Terence Tao is professor of mathematics at the University of California Los Angeles. His e-mail address is tao@math.ucla.edu. *He is supported by the Clay Mathematics Institute and by grants from the Sloan and Packard Foundations.*

Let us denote the eigenvalues of A, B, and $A+B$ as λ, μ, and ν respectively; thus λ_2 is the second largest eigenvalue of A, etc. It is fairly easy to obtain necessary conditions on the triple λ, μ, ν. For instance, from the simple observation that the trace of $A+B$ must equal the sum of the traces of A and B, we obtain the condition

$$(1) \quad \nu_1 + \ldots + \nu_n = \lambda_1 + \ldots + \lambda_n + \mu_1 + \ldots + \mu_n.$$

Another immediate constraint is that

$$(2) \quad \nu_1 \leq \lambda_1 + \mu_1,$$

since the largest eigenvalue of $A+B$ is at most the sum of A's and B's individual largest eigenvalues. (Exercise for the reader: Show equality occurs exactly when the same vector is a principal eigenvector for both matrices.) Weyl found a number of similar necessary conditions, such as the statement $\nu_{i+j+1} \leq \lambda_{i+1} + \mu_{j+1}$ whenever $0 \leq i, j, i+j < n$. When $n = 1, 2$ these conditions are both necessary and sufficient; for higher dimensions many other necessary conditions were found by later authors. All of these conditions took the form of homogeneous linear inequalities (e.g., $\nu_1 + \nu_2 \leq \lambda_1 + \lambda_2 + \mu_1 + \mu_2$). These inequalities were generally proven by "minimax" methods, but there did not appear to be a general scheme that would produce a systematic and complete list of these inequalities.

This problem was studied extensively by Alfred Horn [Ho]. Among other things, he showed that a complete set of necessary conditions could be given by (1), together with a list of linear inequalities of the form

$$(3) \quad \nu_{k_1} + \ldots + \nu_{k_r} \leq \lambda_{i_1} + \ldots + \lambda_{i_r} + \mu_{j_1} + \ldots + \mu_{j_r}$$

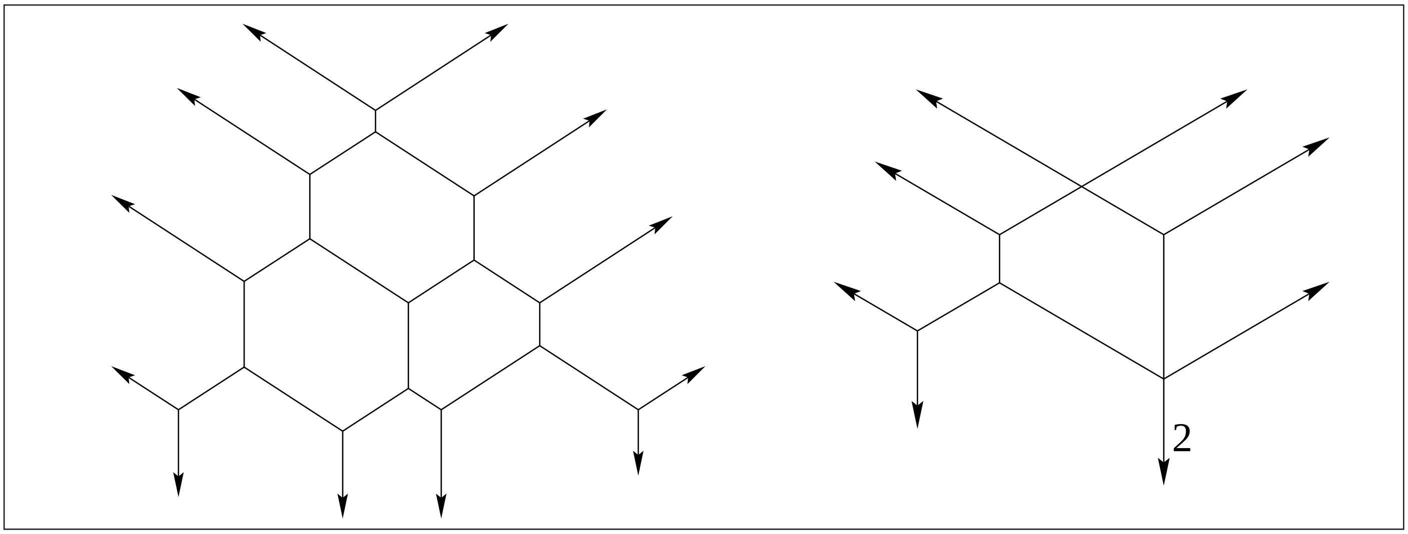

Figure 1. Two honeycombs. The left one is more typical in having only Y vertices, as will be explained in Theorem 3. All edges are multiplicity 1, except for the edge labeled 2 in the right-hand honeycomb.

for all $1 \le r < n$ and all triplets of indices $1 \le i_1 < \ldots < i_r \le n$, $1 \le j_1 < \ldots < j_r \le n$, and $1 \le k_1 < \ldots < k_r \le n$ in a certain finite set $T_{r,n}$. The problem was then reduced to describing the sets $T_{r,n}$ of triplets.

Horn computed this set for $n \le 8$ and for general dimensions was able to demonstrate that the indices $i_1, \ldots, k_r$ in $T_{r,n}$ satisfied the trace condition

$$i_1 + \ldots + i_r + j_1 + \ldots + j_r = k_1 + \ldots + k_r + r(r+1)/2 \tag{4}$$

as well as linear inequalities such as

$$i_1 + j_1 \le k_1 + 1.$$

These relations were obviously similar to the relations (1), (2) in the original problem. This led to the remarkable

Conjecture 1 (Horn conjecture). *The set $T_{r,n}$ is equal to the set of all triplets of indices $1 \le i_1 < \ldots < i_r \le n$, $1 \le j_1 < \ldots < j_r \le n$, $1 \le k_1 < \ldots < k_r \le n$ which obey* (4) *and*

$$i_{a_1} + \ldots + i_{a_s} + j_{b_1} + \ldots + j_{b_s} \ge k_{c_1} + \ldots + k_{c_s} + s(s+1)/2$$

for all $1 \le s < r$ and all triplets of indices $1 \le a_1 < \ldots < a_s \le r$, $1 \le b_1 < \ldots < b_s \le r$, $1 \le c_1 < \ldots < c_s \le r$ in $T_{s,r}$.

This conjecture would give a highly recursive (but impractical) algorithm to generate the sets $T_{r,n}$ in terms of the earlier generations $T_{s,r}$ and thus to give a complete solution to Weyl's problem at each dimension n. The conjecture turns out to be correct, though it waited thirty-six years for resolution.

We approached this problem by first observing that Weyl's problem could be rephrased using *honeycombs*, which we introduced (for this purpose) in [KT]. These are a family of planar arrangements of edges labeled with multiplicities (some examples are in Figure 1). We give the precise definition in the next section.

The relevance of honeycombs to sums of Hermitian matrices is the following theorem, which we explain in more detail later.

Theorem 1. *Let λ, μ, ν be weakly decreasing n-tuples of real numbers. Then there exist matrices A, B, and $A+B$ with respective eigenvalues λ, μ, and ν if and only if there exists a honeycomb with boundary values $(\lambda, \mu, -\nu)$.*

Although the problem about sums of Hermitian matrices is classical, a quantum analogue concerning $U(n)$ representations turns out to be crucial to the resolution of Horn's conjecture. As we shall see, there is also a quantum version of Theorem 1 linking this representation theory problem to (integer) honeycombs. One of the key steps in the proof of Horn's conjecture is the proof of the *saturation conjecture*, which asserts that the classical and quantum problems are in a certain sense equivalent.

We shall give a rather ahistorical (and pro-honeycomb) survey of this circle of ideas, starting with honeycombs (which were actually the last piece of the puzzle to be discovered), then discuss the connections between the classical and quantum problems, followed by a sketch of the honeycomb-based proof of the saturation conjecture. Then we restate Horn's conjecture and sketch the honeycomb-based proof of this from saturation.

There are many other closely related and interesting mathematical questions that we will not address, and we refer the reader to the excellent survey article [F2].

Honeycombs

We now set up some notation needed to define honeycombs and their relation to Weyl's problem.

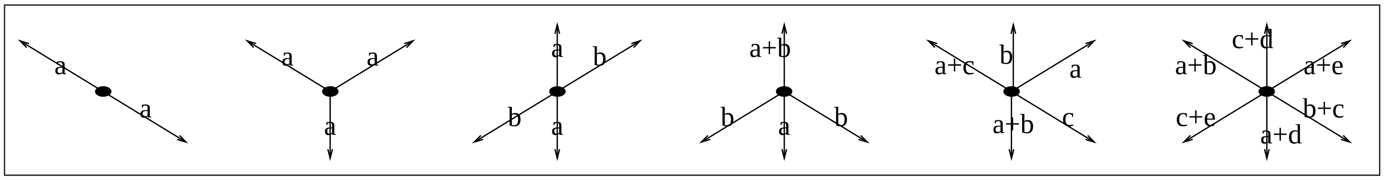

Figure 2. Zero-tension points, with the rays labeled by their multiplicities, which are positive integers. All but the first type are called "vertices".

In the one-dimensional case $n = 1$, a necessary and sufficient set of conditions on λ, μ, ν is given of course by $\lambda + \mu = \nu$. Using the one-dimensional case as an analogy, we then define the relation

$$\lambda \boxplus \mu \sim_c \nu \tag{5}$$

if there exist Hermitian matrices A, B, C with eigenvalues λ, μ, ν respectively such that $A + B = C$. The "c" in $\sim_c$ stands for "classical"; we will define a quantum analogue

$$\lambda \boxplus \mu \sim_q \nu \tag{6}$$

later on. Weyl's problem is thus to determine the solution set to (5).

It is convenient to rephrase Weyl's problem in a more symmetric form. We say that the relation

$$\lambda \boxplus \mu \boxplus \nu \sim_c 0 \tag{7}$$

holds if there exist Hermitian matrices A, B, C with eigenvalues λ, μ, ν respectively such that $A + B + C = 0$. Clearly we have

$$\lambda \boxplus \mu \sim_c \nu \iff \lambda \boxplus \mu \boxplus (-\nu) \sim_c 0$$

where $-\nu := (-\nu_n, \ldots, -\nu_1)$ is the negation of ν. Thus to solve Weyl's problem, it suffices to determine the set of triples λ, μ, ν which obey (7). This formulation has the advantage of S_3 symmetry in (λ, μ, ν), as opposed to mere S_2 symmetry in (λ, μ).

In one dimension $n = 1$, we of course have

$$\lambda \boxplus \mu \boxplus \nu \sim_c 0 \iff \lambda + \mu + \nu = 0.$$

In more general dimensions we have the necessary condition

$$\lambda_1 + \ldots + \lambda_n + \mu_1 + \ldots + \mu_n + \nu_1 + \ldots + \nu_n = 0, \tag{8}$$

which is the analogue of (1). Similarly, (2) becomes

$$\lambda_1 + \mu_1 + \nu_n \geq 0. \tag{9}$$

Based on these relations, it is natural to introduce the plane

$$\mathbb{R}^3_{\sum=0} := \{(x, y, z) \in \mathbb{R}^3 : x + y + z = 0\}.$$

We shall always depict this plane with the six "cardinal directions" $(0, 1, -1), (-1, 1, 0), (-1, 0, 1), (0, -1, 1), (1, -1, 0)$, and $(1, 0, -1)$, drawn northwest, north, northeast, southeast, south, and southwest respectively. (Of course, "northwest" makes a 60° angle with north rather than a 45° angle, and similarly for the other diagonal cardinal directions.)

Define a *diagram* to be a configuration of (possibly half-infinite) line segments in $\mathbb{R}^3_{\sum=0}$, with each edge parallel to one of the cardinal directions (north-south, northeast-southwest, northwest-southeast) and labeled with a positive integer which we refer to as the "multiplicity" or "tension". To every diagram we can associate a measure on $\mathbb{R}^3_{\sum=0}$, defined as the sum of Lebesgue measure on each line segment, weighted by the multiplicity. We say that two diagrams d, d' are *equivalent* if their associated measures are equal.

If h is a diagram and ν is a point in $\mathbb{R}^3_{\sum=0}$, we say that ν is a **zero-tension point** of h if, in a sufficiently small neighborhood of ν, h is equivalent to a union of rays emanating from ν and the sum of the coordinate vectors of these rays, multiplied by their tensions, equals zero.

The two possibilities that will interest us most are a point on a line segment, in which case the zero-tension condition says that the two rays must have the same multiplicity, and a point at the center of a Y with again three equal-multiplicity rays. There are several more complicated possibilities, as shown in Figure 2.

Define a **honeycomb** h as a diagram (or, more precisely, an equivalence class of diagrams) such that

1. every point in $\mathbb{R}^3_{\sum=0}$ is a zero-tension point
2. there are only finitely many "vertices", i.e., points with more than two rays emanating
3. the semi-infinite lines go only in the northeast, northwest, and south directions (i.e., no southeast, southwest, or north rays)

The lines mentioned in number 3 are called the **boundary edges** of the honeycomb. Two examples of honeycombs appear in Figure 1.

It is a pleasant exercise to show that the number of boundary edges (with multiplicity) pointing in one cardinal direction is the same as the number in each of the other two directions. (This is basically because the net tension of the honeycomb must be zero.) We will call a honeycomb with n boundary edges in each direction an n-**honeycomb** and denote the space of such by $\mathtt{HONEY}_n$.

Since every edge in a honeycomb is parallel to one of the cardinal directions, each of which has one of its three coordinates equal to zero, every honeycomb edge has a **constant coordinate** (common

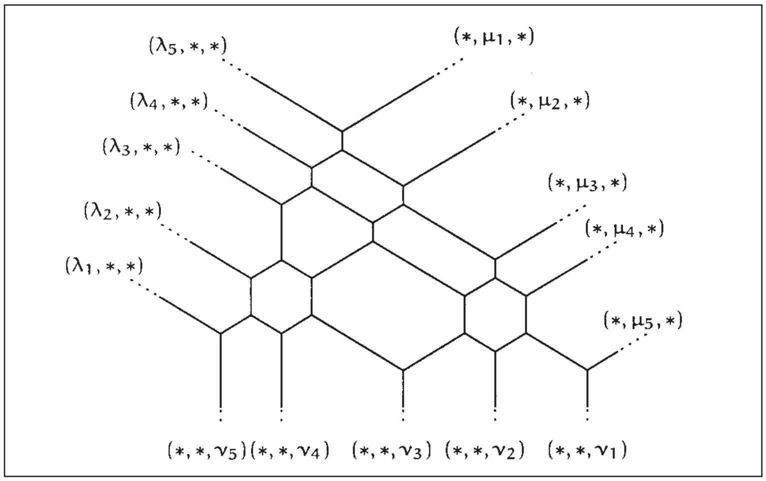

Figure 3. The constant coordinates on the boundary edges of a 5-honeycomb. (The stars are the nonconstant coordinates.)

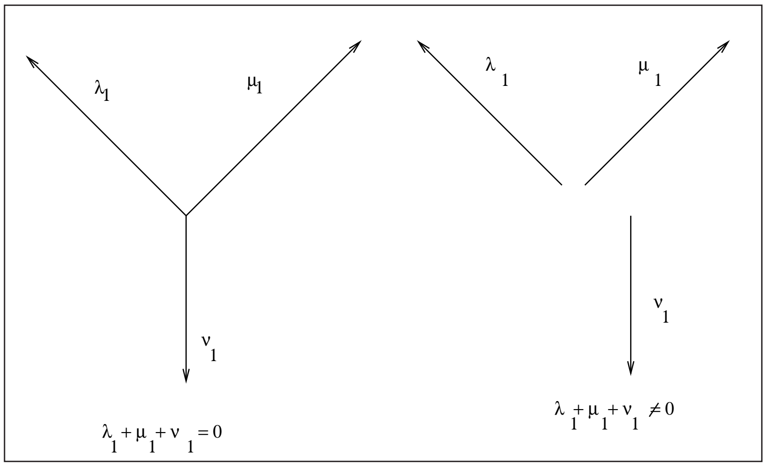

Figure 4. A 1-honeycomb can be formed if and only if the boundary values sum to zero. The edges are labeled by their constant coordinates.

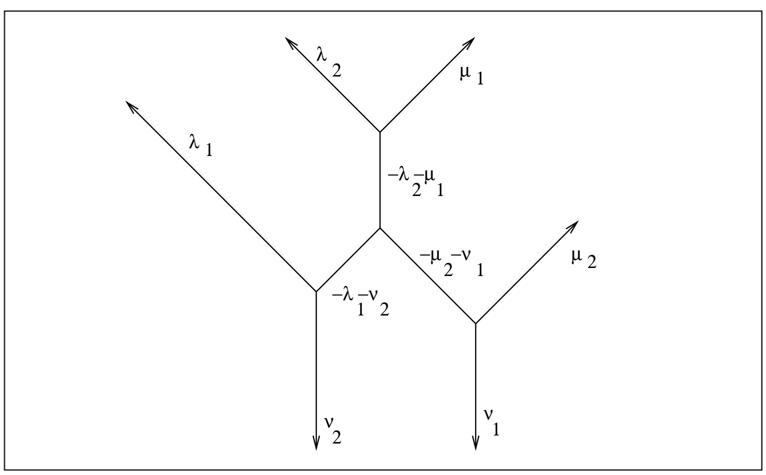

Figure 5. A 2-honeycomb is uniquely determined by its boundary values. The boundary values must satisfy (8), and the edge lengths must be nonnegative. The edges are labeled by their constant coordinates. The edge lengths can be computed (up to an irrelevant factor of $\sqrt{2}$) by subtracting the constant coordinates of two parallel adjacent edges; for instance, the lower left edge has length $\lambda_1 - (-\mu_2 - \nu_1)$.

to every point along that edge). In particular, we can read off the constant coordinates of boundary edges and call them

$$(\lambda_1, \ldots, \lambda_n, \mu_1, \ldots, \mu_n, \nu_1, \ldots, \nu_n) = (\lambda, \mu, \nu)$$

as in Figure 3. (In our pictures of honeycombs we use roman letters to denote multiplicities (which are positive integers), and Greek letters to denote constant coordinates (which are real numbers, but are often integers as well).)

We can now phrase Theorem 1 in this symmetrized setting:

Theorem 2. *The relationship $\lambda \boxplus \mu \boxplus \nu \sim_c 0$ holds if and only if there exists a honeycomb with boundary values (λ, μ, ν).*

Interestingly, almost all the proofs we know of this theorem proceed by first proving a quantized version, which we define in a later section. We shall therefore not discuss the proof of this theorem here, and content ourselves instead with producing evidence which strongly suggests that the theorem is plausible.

We first consider the $n = 1$ case. In this case $\lambda = (\lambda_1)$, $\mu = (\mu_1)$, $\nu = (\nu_1)$, and it is clear that (7) holds if and only if $\lambda_1 + \mu_1 + \nu_1 = 0$. (In other words, the trace condition is already necessary and sufficient.) On the honeycomb side this claim can be easily seen if one accepts the fact (which is actually a little tricky to prove) that 1-honeycombs must have the shape of a "Y". See Figure 4.

More generally, it is a pleasant exercise to show that the boundary values of any n-honeycomb must satisfy (8), basically because the three coordinates around every vertex sum to zero (by virtue of lying in $\mathbb{R}^3_{\sum=0}$).

Now consider the $n = 2$ case, so that $\lambda = (\lambda_1, \lambda_2)$, $\mu = (\mu_1, \mu_2)$, $\nu = (\nu_1, \nu_2)$. In this case there can be at most one 2-honeycomb with the specified boundary values (Figure 5). The lengths of the three line segments in the honeycomb can be computed as $\lambda_2 + \mu_1 + \nu_1$, $\lambda_1 + \mu_2 + \nu_1$, $\lambda_1 + \mu_1 + \nu_2$. Since these line segments need to have nonnegative length, we obtain the necessary conditions

$$\lambda_2 + \mu_1 + \nu_1, \lambda_1 + \mu_2 + \nu_1, \lambda_1 + \mu_1 + \nu_2 \geq 0.$$

These inequalities can be rephrased using (8) as the statement that the quantities $\lambda_1 - \lambda_2$, $\mu_1 - \mu_2$, $\nu_1 - \nu_2$ form the side-lengths of a triangle. The reader may verify from some linear algebra that these conditions are indeed necessary and sufficient for (7).

In the $n > 2$ case things become more complicated, because the boundary values no longer uniquely determine the honeycomb. In fact, every hexagon present in a honeycomb provides a degree of freedom; the hexagon can be "breathed" inwards or outwards (see Figure 6).

However, it is still possible to demonstrate that inequalities such as (9) must hold for n-honeycombs. Indeed, one can simply extend the μ_1 ray southward until it intersects the λ_1

ray. This intersection point must be northwest of the intersection of ν_n and λ_1, which gives (9). More generally, the Weyl inequalities $\lambda_i + \mu_j + \nu_k \geq 0$ for $i + j + k = n + 2$ can be demonstrated by constructing a Y-shaped object embedded inside the n-honeycomb which is quite similar to a 1-honeycomb (see Figure 7). Similarly, inequalities involving pairs of eigenvalues can be demonstrated by constructing an object similar to a 2-honeycomb; the reader may be amused by locating the object needed to prove $\lambda_1 + \lambda_4 + \mu_1 + \mu_4 + \nu_{n-4} + \nu_{n-1} \geq 0$. A more careful pursuit of this idea can be used to obtain half of Horn's conjecture (that lower-order honeycombs generate inequalities for higher-order honeycombs). The other half, that all inequalities for honeycombs are generated in this way, is proven by the machinery of transverse clockwise overlays, which we discuss later.

Having given some examples of how necessary conditions for (7) translate to the honeycomb setting, we now look at sufficient conditions. It is easy to see (by restricting A, B, C to diagonal matrices) that (7) will hold if there exist permutations $\alpha, \beta \in S_n$ such that $\lambda_{\alpha(i)} + \mu_{\beta(i)} + \nu_i = 0$ for all $1 \leq i \leq n$. The honeycomb analogue of this is depicted in Figure 8; one can obtain a (rather degenerate) n-honeycomb by overlaying n separate 1-honeycombs on top of one another.

More generally, there is a notion of *overlaying* an n-honeycomb h and an m-honeycomb h' to form an $(n+m)$-honeycomb $h \oplus h'$. To be precise, $h \oplus h'$ is the honeycomb whose associated measure is the sum of the measures associated to h and h'. This operation corresponds to the direct sum operation on Hermitian matrices (which takes an $n \times n$ matrix and an $m \times m$ matrix and forms an $(n+m) \times (n+m)$ block-diagonal matrix) or the concatenation operation on spectra (which takes a set of n eigenvalues and a set of m eigenvalues and forms the (sorted) set of $n+m$ eigenvalues). Intuitively, an overlay can be demonstrated by drawing two honeycombs on transparencies and stacking both transparencies on the same projector; see Figure 9. We shall have more to say about overlays later in this article.

The statement that (λ, μ, ν) admits a honeycomb with these boundary values is clearly symmetric under cyclic permutations. However, the relation (7) is symmetric under the larger group S_3, thanks to the commutativity of addition. The corresponding S_3 symmetry result for honeycombs is not trivial; an elegant proof of this based on scattering arguments is given in [Wo]. The same argument also gives the associativity property

$$(10) \quad (\exists \nu : \lambda \boxplus \mu \sim_c \nu; \nu \boxplus \rho \sim_c \sigma) \iff (\exists \nu' : \mu \boxplus \rho \sim_c \nu'; \lambda \boxplus \nu' \sim_c \sigma);$$

this property, combined with a "Pieri rule" to handle the generating cases when λ or μ is equal to

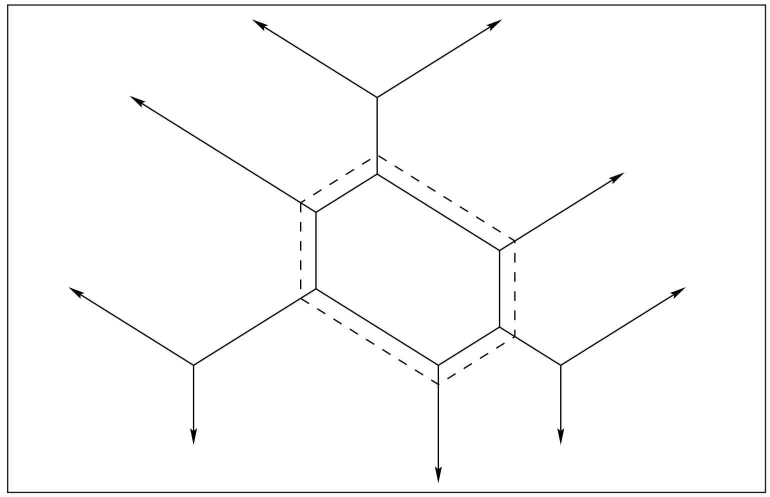

Figure 6. A hexagon in a honeycomb, with a dotted line indicating a place to which one might dilate it.

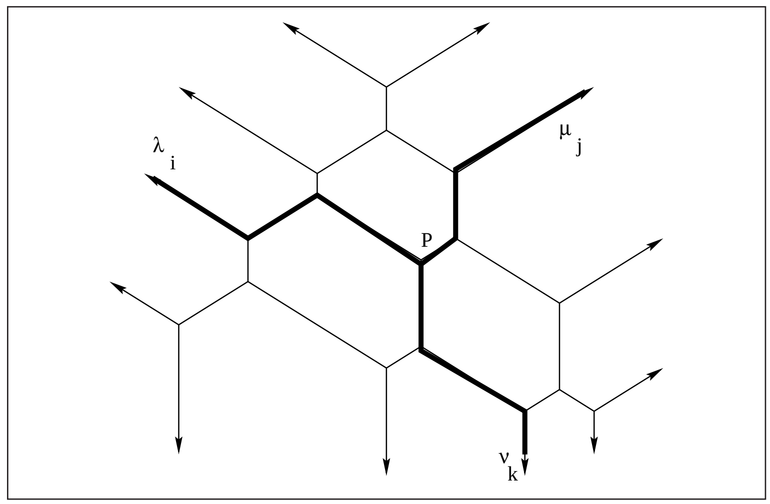

Figure 7. A honeycomb proof of an inequality $\lambda_i + \mu_j + \nu_k \geq 0$. The three coordinates of point P must sum to zero; however, the first coordinate cannot exceed λ_i, the second cannot exceed μ_j, and the third cannot exceed ν_k, hence the claim. Note the resemblance between the shape drawn in bold and a 1-honeycomb.

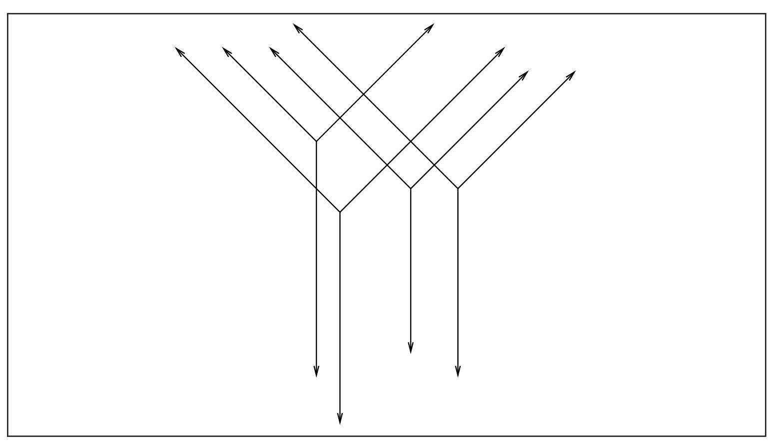

Figure 8. An n-honeycomb can be obtained by overlaying n 1-honeycombs on top of one another. Here α and β map $1, 2, 3, 4$ to $4, 3, 1, 2$ and $4, 3, 2, 1$ respectively.

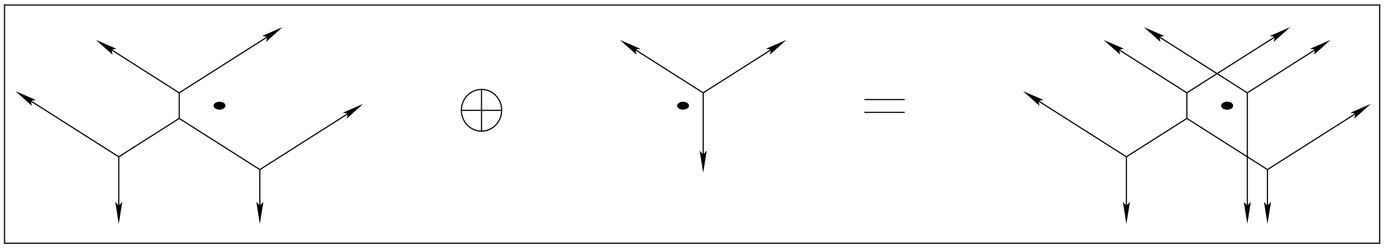

Figure 9. Two honeycombs overlaid to produce a third. The origin $(0,0,0)$ is marked in each picture with a black dot.

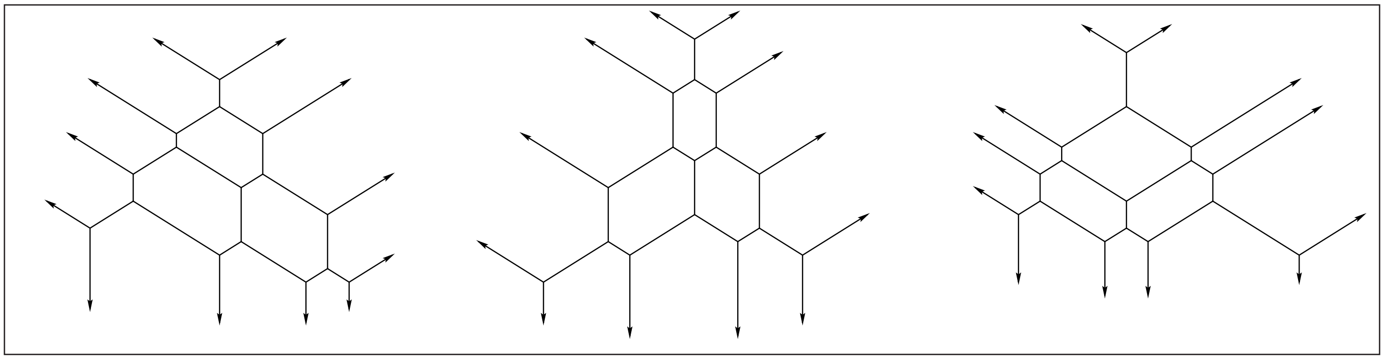

Figure 10. Three nondegenerate 4-honeycombs. Note that there is a natural way to correspond the edges in one with the edges in any other.

$(\varepsilon, 0, \ldots, 0)$ for some small ε, can be used to give an inductive proof of Theorem 1.

An interesting degenerate case occurs when ν is kept fixed, while the spacings between eigenvalues of λ and μ are allowed to become very large. In this case the honeycomb degenerates into a pattern known as a Gelfand-Cetlin pyramid, while Weyl's problem degenerates to Schur's problem of determining which n-tuples can be the diagonal entries of a Hermitian matrix with specified eigenvalues. (In fact, we discovered honeycombs by extrapolating from this degenerate case.)

We encourage the reader to try out the honeycomb Java applet at `http://www.math.ucla.edu/~tao/java/Honeycomb.html`.

Organizing Honeycombs into a Polyhedral Cone

The space $\mathtt{HONEY}_n$ of all n-honeycombs has been defined as an abstract set, but one can, in fact, give this space the structure of a polyhedral cone inside some finite-dimensional vector space.

Call a honeycomb **nondegenerate** if

1. all its edges are multiplicity 1, and
2. all its vertices are either right-side-up or upside-down Ys.

It is straightforward to prove that all nondegenerate n-honeycombs have the same topological shape, namely that of Figure 10. In particular, there is a natural one-to-one correspondence between the edges in one nondegenerate n-honeycomb with those in any other.

This gives us a way of making the space of nondegenerate n-honeycombs into an open polyhedral cone in a real vector space. The coordinates are given by the constant coordinates of the edges, and linear constraints are imposed by saying that the constant coordinates of three edges meeting at a vertex add to zero and that every edge has strictly positive length.

Theorem 3. [KT] *The identification in the above paragraph between nondegenerate n-honeycombs and points in a certain polyhedral open cone extends to an identification of all $\mathtt{HONEY}_n$ with the closure of this cone. In particular, nondegenerate n-honeycombs form a dense open set in $\mathtt{HONEY}_n$.*

This theorem is surprisingly annoying to prove and takes pp. 1067–1074 of [KT]. Its virtue is in enabling us to use the theory of such cones to define certain special honeycombs (the "largest lift" honeycombs). The proof of the saturation conjecture hinges on the fact that every honeycomb can be deformed to a largest lift honeycomb.

For a honeycomb h, let $\partial h \in (\mathbb{R}^n)^3$ denote the list $(\lambda_1, \ldots, \lambda_n, \mu_1, \ldots, \mu_n, \nu_1, \ldots, \nu_n)$ of constant coordinates on the boundary edges of h, and let $\mathtt{BDRY}_n \subset (\mathbb{R}^n)^3$ be the image of this map $\partial : \mathtt{HONEY}_n \to (\mathbb{R}^n)^3$. Then we can think of the main question as being to list the inequalities determining $\mathtt{BDRY}_n$.

It is not hard to show directly that this map is proper, so each fiber is a compact, convex polyhedron.

As one application of this formalism we can easily show that for any λ, μ, ν the truth or falsity of (7) or (5) can be determined in polynomial time with respect to the dimension n; this fact appears to be previously unknown. Indeed, the problem is equivalent to determining whether the polytope $\partial^{-1}(\lambda, \mu, \nu)$ is nonempty. Since the cone $\mathtt{HONEY}_n$ has about $O(n^2)$ faces, which can all be described explicitly, this problem can be decided in polynomial time by standard algorithms (e.g., the simplex method). [The authors thank Peter Shor for pointing out this fact.] On the other hand, we do not know how to enumerate all the determining inequalities for the relationship (7) in an efficient manner; the recipe given by Horn's conjecture requires worse-than-exponential time and memory in n and in fact produces many redundant inequalities for (7).

One can also use this formalism to create a more quantitative version of Theorems 1 and 2. Let $\mathcal{O}_\lambda$ denote the manifold of Hermitian matrices with eigenvalues λ, and let A be the random variable with the uniform distribution on $\mathcal{O}_\lambda$ (where "uniform" can be defined using induced Lebesgue measure, or the $U(n)$ action). In other words, A is a random matrix with spectrum λ. Similarly, define B as a random matrix with spectrum μ. One can then define $P(\lambda \boxplus \mu \sim_c \nu)$ to be the probability density of the spectrum of the sum $A + B$ of two independent random matrices evaluated at ν. Similarly, define $P(\lambda \boxplus \mu \boxplus \nu \sim_c 0)$.

Theorem 4. *Up to inessential factors (constants and Vandermonde determinants),* $P(\lambda \boxplus \mu \sim_c -\nu)$ *and* $P(\lambda \boxplus \mu \boxplus \nu \sim_c 0)$ *are equal to the volume of* $\partial^{-1}(\lambda, \mu, \nu)$*.*

Readers familiar with symplectic geometry will recognize this type of theorem from the theory of moment maps of compact Lie groups such as $U(n)$. Indeed, $P(\lambda \boxplus \mu \sim_c -\nu)$ is essentially the volume of the symplectic reduction of the manifold $\mathcal{O}_\lambda \times \mathcal{O}_\mu$ (with the diagonal $U(n)$ action) at the point $-\nu$ and similarly for $P(\lambda \boxplus \mu \boxplus \nu \sim_c 0)$. We shall have more to say about this later on.

Quantum Analogues

We now describe the quantum analogue (6) of the classical relation (5). Roughly speaking, (6) is to the representation theory of $U(n)$ as (5) is to the symplectic geometry of $U(n)$ (or, more precisely, of the coadjoint orbits $\mathcal{O}_\lambda$ of $U(n)$).

Recall that the irreducible unitary representations of $U(1)$ are all one-dimensional. In fact, for each integer λ we can define the irreducible representation V_λ as a one-dimensional vector space, with the action of $e^{i\theta}$ given by multiplication by $e^{i\lambda\theta}$ on V_λ.

More generally, for any weakly decreasing sequence $\lambda = (\lambda_1 \geq \ldots \geq \lambda_n)$ of integers we can define an irreducible unitary representation V_λ of $U(n)$ by standard constructions (see, e.g., [F]). The n-tuple λ is referred to as the *weight* of V_λ. For instance, if λ consists of k ones and $n-k$ zeroes, then V_λ is the space of k-forms $\bigwedge^k \mathbb{C}^n$ with the standard $U(n)$ action. More generally, if $\lambda_n \geq 0$, we define V_λ to be the highest-weight irreducible representation in

$$\bigotimes_{i=1}^{n} Sym^{\lambda_i - \lambda_{i+1}} \bigwedge^i \mathbb{C}^n$$

with the convention $\lambda_{n+1} = 0$, and the $\lambda_n < 0$ representations can be defined via a dualization.

Given two irreducible representations V_λ, V_μ of $U(n)$, the tensor product $V_\lambda \otimes V_\mu$ is another representation of $U(n)$. In the $n = 1$ case the tensor product is again an irreducible representation: $V_\lambda \otimes V_\mu \equiv V_{\lambda+\mu}$. However, in general the tensor product is not irreducible and splits up as a direct sum of many smaller irreducible representations V_ν. We can now define the relation (6) as the statement that a copy of V_ν appears at least once in the tensor product $V_\lambda \otimes V_\mu$. Note that the quantum relation is only defined for *integral* λ, μ, ν, whereas the classical relation (5) is defined for *real* λ, μ, ν.

There is a close parallel between (5) and (6). For instance, one can obtain the trace identity (1) as a necessary condition for (6) by considering the action of the center $U(1)$ of $U(n)$. One can similarly obtain the necessary condition (2) by considering the highest weights of the action of a maximal torus $U(1) \times \ldots \times U(1)$ in $U(n)$. From a more physical viewpoint, one can view the classical problem as a problem of describing how the moments of inertia of bodies in $\mathbb{C}^n$ behave under superposition, while the quantum problem is the problem of describing how the spin states of particles in $\mathbb{C}^n$ behave under superposition. (The $n = 2$ case is especially interesting to physicists, since $U(2)$ is closely related to $O(3)$. In this case every representation V_ν appears at most once in $V_\lambda \otimes V_\mu$ (this corresponds to the fact that 2-honeycombs are determined by their boundary values), and one can parameterize the decomposition explicitly using the Clebsch-Gordan coefficients.)

This connection between the classical and quantum problems seems to have been noted first in [L] (and in a more general context in [He], both in 1982) and appears in detail in [Kl]; the most natural framework for such results is exposed in [Kn]. Explicitly, the connection is given by

Theorem 5. *Let* λ, μ, ν *be weakly decreasing sequences of n integers.*

1. *(Quantum implies classical.) If (6) holds, then (5) holds.*
2. *(Classical implies asymptotic quantum.) Conversely, if (5) holds, then there exists an integer* $N > 0$ *such that* $N\lambda + N\mu \sim_q N\nu$*. (Here* $N\lambda$ *is the sequence* $(N\lambda_1, \ldots, N\lambda_n)$*.)*

From this theorem it is natural to phrase

Conjecture 2 (saturation conjecture). *One can take* $N = 1$ *in the above theorem. In other words, (5) and (6) are equivalent for integer* λ, μ, ν*.*

This conjecture seems to be special to $U(n)$; the naïve analogue of this conjecture for other Lie groups can be easily shown to be false. The saturation conjecture is so named because it is equivalent to the set of triples (λ, μ, ν) obeying (6) being a saturated submonoid of $\mathbb{Z}^{3n}$.

Using some formidable algebraic and geometric machinery, Klyachko [Kl] was able to demonstrate a further nontrivial recursive relationship between the classical and quantum problems and noted that this, combined with the saturation conjecture, would imply Horn's conjecture; we shall have more to say about this later. In [KT] we used Theorem 1 (and the quantum analogue of this theorem) to convert the saturation conjecture into a statement about honeycombs and then proved this statement by combinatorial methods, thus proving the saturation and Horn conjectures. (Recently it has been shown [KTW] that one can derive Horn's conjecture directly from the saturation conjecture by purely combinatorial techniques, bypassing the machinery of [Kl]. Also, a very different proof of saturation, based on the representation theory of quivers, has since been given in [DW]. Finally, a short rendition of [KT] can be found in [Bu].)

To attack the saturation conjecture using honeycombs, we need a quantum analogue of Theorems 1 and 2. We first phrase a symmetric form of (6). We say that

$$\lambda \boxplus \mu \boxplus \nu \sim_q 0 \tag{11}$$

holds if $V_\lambda \otimes V_\mu \otimes V_\nu$ contains a nontrivial $U(n)$-invariant vector. It is easy to show that (11) is equivalent to $\lambda \boxplus \mu \sim_q -\nu$.

A honeycomb is said to be *integral* if its vertices lie on $\mathcal{Z}^3_{\sum=0} := \mathbb{R}^3_{\sum=0} \cap \mathbb{Z}^3$. Note that the boundary values of an integral honeycomb are necessarily integers.

Theorem 6. *The relationship (11) holds if and only if there exists an integral honeycomb with boundary values (λ, μ, ν). As a corollary the relationship (6) holds if and only if there exists an integer honeycomb with boundary values $(\lambda, \mu, -\nu)$.*

Note that Theorems 6 and 5 imply Theorems 1 and 2.

The problem of determining the solutions to (6) has had a long history, and a solution is given by the famous Littlewood-Richardson rule. This rule has been formulated in many different ways, most of which involve Young tableaux; a variant due to Berenstein and Zelevinsky can be easily adapted to give Theorem 6. (Fulton has also shown that this theorem can be proven directly from the Littlewood-Richardson rule.) Other proofs are known; for instance, one can combine the quantum version of (10) with Pieri's rule for tensoring a $U(n)$ representation with the tautological $\mathbb{C}^n$ representation to give an inductive proof of Theorem 6.

A quantum analogue of Theorem 4 is also known:

Theorem 7. *The number of times V_ν appears in the tensor product of $V_\lambda \otimes V_\mu$ is equal to the number of integral honeycombs with boundary values $(\lambda, \mu, -\nu)$. Equivalently, the dimension of the $U(n)$-invariant subspace of $V_\lambda \otimes V_\mu \otimes V_\nu$ is equal to the number of integral honeycombs with boundary values (λ, μ, ν).*

All the proofs of Theorem 6 mentioned above can also be used to prove Theorem 7. Theorem 4 can be viewed as a crude asymptotic version of Theorem 7. Variants of this theorem appear in [J], [BZ], and particularly in [GP], though honeycombs are not explicitly used in these papers. We remark that the representation theoretic quantities in Theorem 7 can also be calculated by the Steinberg product rule (for instance), though we do not know a proof of this theorem that goes via this rule.

Readers who are familiar with the representation theory of SL_2 (or $SU(2)$) may verify that the honeycomb rule given in Theorem 6 corresponds to the usual triangle inequalities for the weights. The fact that 2-honeycombs are uniquely determined by their boundary values corresponds to the fact that each irreducible representation of SL_2 appears exactly once in a tensor product of irreducibles.

As mentioned in the introduction, the saturation conjecture gives a complete solution to the (now equivalent) problems (5), (6), given by Horn's conjecture.

Proof of the Saturation Conjecture

In light of the theorems of the previous section, the saturation conjecture can be reduced to the following purely honeycomb-theoretic problem:

Theorem 8. *Let h be a (real-valued) honeycomb with integer boundary values. Then there exists an integer honeycomb h' with the same boundary values as h.*

Or in other words: if λ, μ, ν are integers and the polytope $\partial^{-1}(\lambda, \mu, \nu)$ is nonempty, then $\partial^{-1}(\lambda, \mu, \nu)$ must contain at least one integer point.

The most obvious thing to do is to look for a vertex of $\partial^{-1}(\lambda, \mu, \nu)$; however, one can give examples of vertices which are nonintegral even when λ, μ, ν are integers. Thus we have to be a little more careful as to how to locate our integer honeycomb.

We call a functional $f : \mathtt{HONEY}_n \to \mathbb{R}$ **superharmonic** if it increases when we dilate a hexagon (which one can do to any hexagon in any nondegenerate honeycomb, as in Figure 6).

Fix a generic superharmonic functional f, and define the **largest lift** of a triple (λ, μ, ν) as the honeycomb h that maximizes $f(h)$ subject to $\partial h = (\lambda, \mu, \nu)$. It is straightforward to prove that the largest-lift map $\mathtt{BDRY}_n \to \mathtt{HONEY}_n$ is uniquely defined (for a given generic f), continuous, and

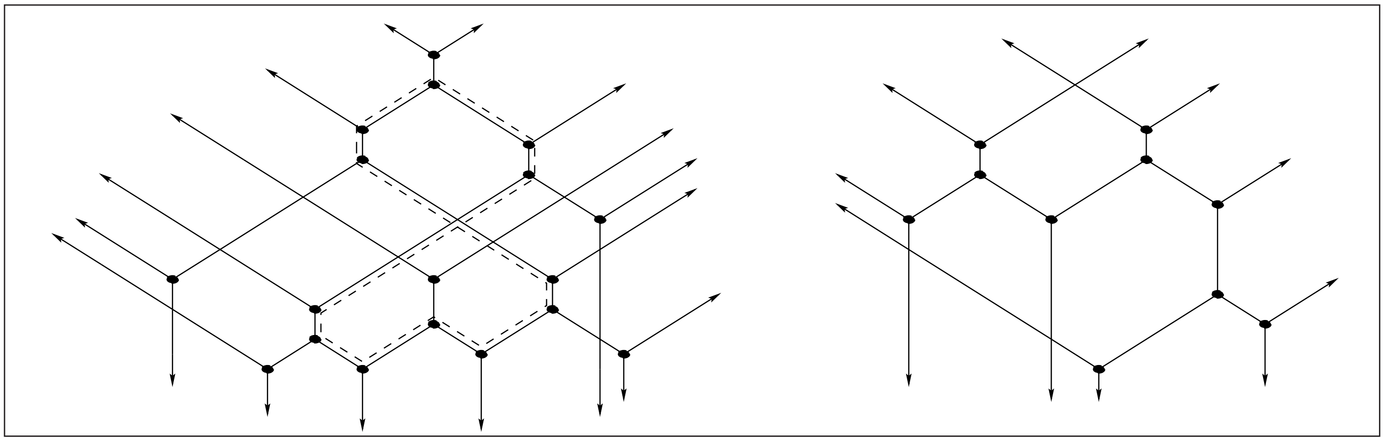

Figure 11. Two simply degenerate honeycombs, with black dots on the vertices of their underlying graphs. The left one has a loop that can breathe in and out (to, say, the dotted-line position), but the right has none.

piecewise-linear. To show Theorem 8, it then suffices to show that the largest-lift map takes integer boundary values to integer honeycombs.

We need some more notation. Say that a honeycomb h has **only simple degeneracies** if all its edges are multiplicity 1 and its vertices are either Ys (possibly upside-down) or crossings of two straight lines. In this case define the **underlying graph of** h as the graph whose vertices consist of the (possibly upside-down) Ys but *not* the crossings; the crossings we instead interpret as two edges missing one another. (In the examples in Figure 1, all vertices *except* the bottom right vertex of the right-hand honeycomb are only simple degeneracies.)

With this in mind, we can talk about **loops** in a simply degenerate honeycomb (meaning in the underlying graph) or call the honeycomb **acyclic** if there are none. For example, in the left honeycomb in Figure 11 there is a loop, whereas the honeycomb on the right is acyclic.

The importance of loops in simply degenerate honeycombs is that they can be breathed in and out, as in Figure 11, generalizing the case of dilating a single hexagon.

Call a largest lift *regular* if the boundary spectra λ, μ, ν each contain no repeated eigenvalues.

The main technical part of [KT] is to prove

Theorem 9. [KT] *Regular largest lifts can only have simple degeneracies.*

In particular, regular largest lifts come with underlying graphs. Roughly speaking, this theorem is proven by showing that every nonsimple degeneracy can be "blown up" in a way that increases the superharmonic functional.

Lemma 1. *The underlying graphs of regular largest lifts are acyclic.*

Sketch of proof. If a simply degenerate honeycomb has a loop, we can breathe it in and out; one direction will increase the value of any (generic) superharmonic functional. A largest lift is by

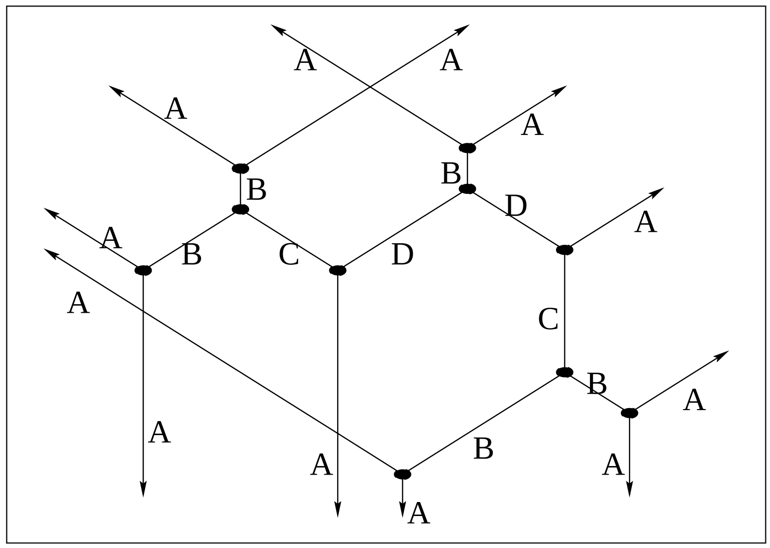

Figure 12. A honeycomb integrally determined by its boundary, in four stages: A, B, C, D.

assumption already at the maximum value of the functional, so there can be no loops. □

From this lemma one can show that the coordinates of regular largest lifts are integral linear combinations of the boundary values. Those who wish to see the details should go to [KT], but the argument is intuitively clear. Given a honeycomb with some edges labeled by their constant coordinates and some still mysterious, look for vertices with two known constant coordinates. The remaining one is minus the sum of the other two. Label it such and repeat. The reader is invited to play this game on the honeycombs in Figure 11 to see how in the left-hand honeycomb one gets stuck exactly because of the loop.

In Figure 12 we have labeled the boundary edges "A", the edges whose constant coordinates can be determined from those "B", those at the next stage of this recursive algorithm "C", and so on.

Since every largest lift can be obtained as a limit of regular largest lifts, we thus have that the

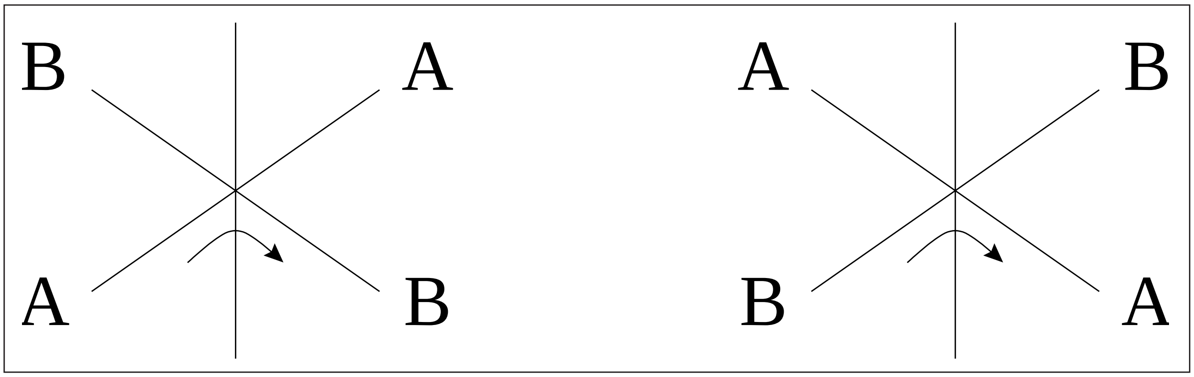

Figure 13. In the left figure, A turns clockwise to B, whereas in the right the reverse is true. Any transverse point of intersection of two overlaid honeycombs must look like exactly one of these.

co-ordinates of all largest lifts are integral linear combinations of their boundary values. In particular, if the boundary values are integral, then the largest lifts are also integral. This proves Theorem 8, which gives the saturation conjecture.

Klyachko's Result and Horn's Conjecture

In this section we restate Horn's conjecture in a convenient form and state a version of Klyachko's result (one direction of which was also proven by Helmke and Rosenthal). (We give slightly revisionist versions in order to avoid introducing Schubert calculus on Grassmannians, which is one of the equivalent problems explained in [F2].)

Horn [Ho] showed that the solution set to (5) must be given by (1) and a finite number of inequalities of the form

$$\lambda_{i_1+r} + \ldots + \lambda_{i_r+1} + \mu_{j_1+r} + \ldots + \mu_{j_r+1} \geq \nu_{k_1+r} + \ldots + \nu_{k_r+1} \tag{12}$$

where $1 \leq r < n$, and $i = (i_1 \geq \ldots \geq i_r)$, $j = (j_1 \geq \ldots \geq j_r)$, and $k = (k_1 \geq \ldots \geq k_r)$ are weakly decreasing sequences of integers between 0 and $n - r$ inclusive. Let us call triples (i, j, k) of this form *admissible*.

As an example, (2) is (12) for the admissible triple $((0), (0), (0))$, while Weyl's inequalities correspond to admissible triples of the form $((i), (j), (i + j))$. The inequality $\lambda_1 + \lambda_2 + \mu_1 + \mu_2 \geq \nu_1 + \nu_2$ corresponds to $((0, 0), (0, 0), (0, 0))$ and so forth.

Horn's conjecture can be easily shown by induction to be equivalent to

Conjecture 3. *Let λ, μ, ν be weakly decreasing sequences of real numbers. Then (5) holds if and only if (1) holds, and (12) holds whenever i, j, k are admissible and $i \boxplus j \sim_c k$.*

Helmke, Rosenthal, and Klyachko showed that Horn's conjecture was true provided that the $\sim_c$ relation on (i, j, k) was replaced by the quantum counterpart $\sim_q$:

Theorem 10. *Let λ, μ, ν be weakly decreasing sequences of real numbers obeying (1).*

1. [HR], [K1] *If (5) holds, then (12) holds whenever (i, j, k) are admissible triples obeying $i \boxplus j \sim_q k$.*
2. [Kl] *Conversely, if (12) holds whenever (i, j, k) are admissible triples obeying $i \boxplus j \sim_q k$, then (5) holds.*

So in a sense solvability of the "classical problem in dimension n" (about summing $n \times n$ Hermitian matrices) is determined by the solvability of the "quantum problem in dimension $m < n$" (about tensoring representations of $U(m)$). Given the saturation theorem proven in the last section, which says that each such quantum problem is solvable exactly if the corresponding classical problem (in the same dimension) is solvable, we have a recursive way to answer the problem.

Theorem 10 connects the classical and quantum problems in a way markedly different from the standard classical/quantum analogy as codified by Theorem 5. The proofs of this theorem are highly nontrivial and first proceed by showing (6) is equivalent to a certain intersection problem in the Schubert calculus of Grassmanians. We do not discuss this further here, but refer the interested reader to [F2]. More recently, a purely honeycomb-theoretic proof of Theorem 10 has been obtained, which we discuss briefly in the last section.

Other Consequences

We close with mention of a few other applications of honeycombs and their properties proven above.

Horn's proof that the solution set of (5) is determined by (1) and a finite number of inequalities of the form (12) is based on the following stronger fact: if (12) holds with equality and λ, μ, ν are regular, then the associated triple of matrices $(A, B, A + B)$ is necessarily block diagonalizable. Put another way, the Hermitian triple is the direct sum of two smaller Hermitian triples.

Given that we have already drawn an analogy between direct sums of matrices and overlaying of honeycombs, there should be a corresponding statement stating that the faces of BDRY_n correspond to honeycombs which are overlays.

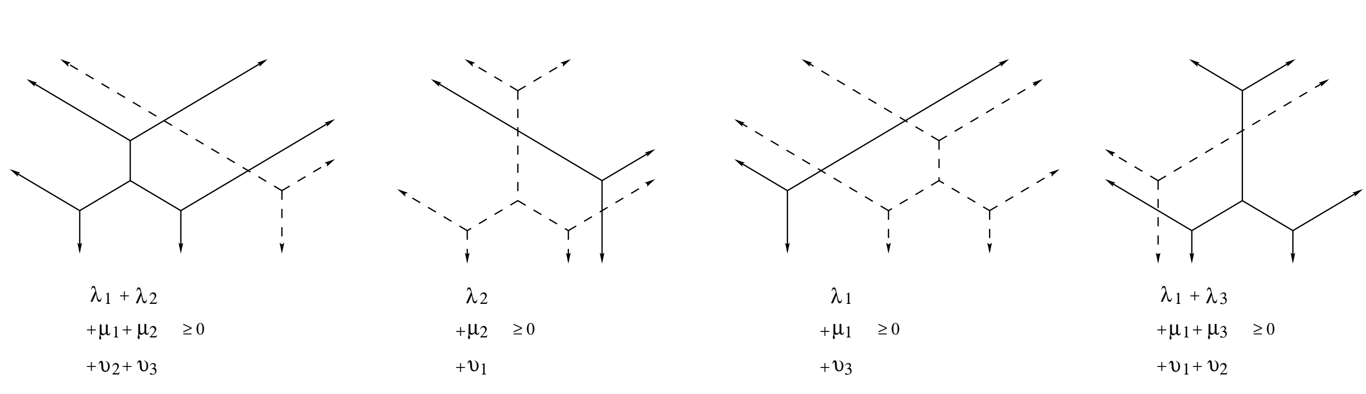

Figure 14. All the consistently clockwise overlays of size 3, up to rotation and deformation, and the associated inequalities on $\mathtt{BDRY}_3$. In each one the solid honeycomb turns clockwise to the dashed honeycomb.

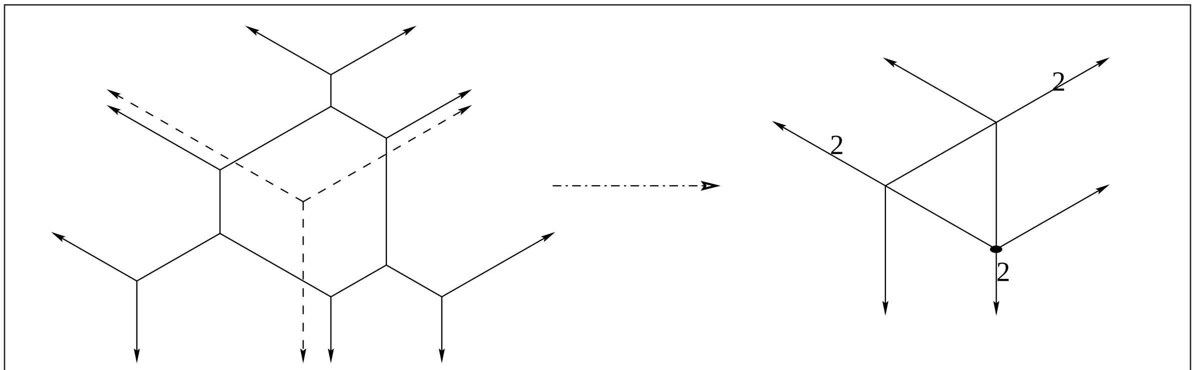

Figure 15. The honeycomb on the right comes from the solid honeycomb on the left, with each edge rescaled to the number of times it intersects the dashed honeycomb and translated to put the bottom right vertex at the origin (marked with a dot).

Not every direct sum of Hermitian matrices (or an overlay of two honeycombs) corresponds to an inequality (12) in Horn's list. However, we have

Theorem 11. [KTW] *Let h be a generic point in $\mathtt{HONEY}_n$. Then $\partial(h)$ is on a facet of $\mathtt{BDRY}_n$ if and only if h is an overlay of two smaller honeycombs A and B, such that at each point of intersection some neighborhood of h looks like the left figure in Figure 13 ("A turns clockwise to B").*

In this case one can read off the inequality on $\mathtt{BDRY}_n$ directly from h: it says that the sum of the boundary coordinates of A is nonnegative. Some examples are in Figure 14.

If A always turns clockwise to B at the intersections, we can construct a new m-honeycomb A' by shifting A so the bottom right vertex is at $(0, 0, 0)$, replacing each of the edges of A by one whose length is the number of intersections with B, and removing B altogether. The result is in fact a new honeycomb, integral and of size m (an example is in Figure 15). This construction can be used to give a purely honeycomb-theoretic proof [KTW] of the results of Klyachko and Helmke-Rosenthal and gives enough additional insight to cut down Horn's overcomplete list of inequalities to the minimum possible.

An Open Question

The present proof of Theorem 4 is very unsatisfying; it comes as an asymptotic limit of Theorem 7, which itself is proved only indirectly.

Consider the horizontal projection of the 2-sphere of height $1/(2\pi)$ onto the diameter between the poles. Archimedes' theorem states that the length of an interval in that diameter equals the area of the preimage on the sphere. Today we say that the horizontal projection is *measure-preserving*, which at first seems marvelous, since the interval has only half the dimension of the sphere. The question is: is there a corresponding map which would give a direct proof of Theorem 4? In other words, is there a canonical measure-preserving map from the set $\{(A, B, C) \in \mathcal{O}_\lambda \times \mathcal{O}_\mu \times \mathcal{O}_\nu\}$ to $\partial^{-1}(\lambda, \mu, \nu)$? Such a map is also likely to give a direct proof of Theorem 7, especially if it is associated somehow with a group action.

References

[Bu] A. BUCH, The saturation conjecture (after A. Knutson and T. Tao), *Enseign. Math.*, to appear.

[BZ] A. BERENSTEIN and A. ZELEVINSKY, Triple multiplicities for $sl(r+1)$ and the spectrum of the exterior algebra of the adjoint representation, *J. Algebraic Combin.* **1** (1992), 7–22.

[DW] H. Derksen and J. Weyman, Semi-invariants of quivers and saturation for Littlewood-Richardson coefficients, *J. Amer. Math. Soc.*, to appear.

[F] W. Fulton, *Young Tableaux*, Cambridge Univ. Press, 1997.

[F2] W. Fulton, Eigenvalues, invariant factors, highest weights, and Schubert calculus, *Bull. Amer. Math. Soc. (N.S.)* **37** (2000), 209–249.

[GP] O. Gleizer and A. Postnikov, Littlewood-Richardson coefficients via Yang-Baxter equation, `math.QA/9909124`, to appear in *Internat. Math. Res. Notices.*

[He] G. J. Heckman, Projections of orbits and asymptotic behavior of multiplicities for compact connected Lie groups, *Invent. Math.* **67** (1982), 333–356.

[HR] U. Helmke and J. Rosenthal, Eigenvalue inequalities and Schubert calculus, *Math. Nachr.* **171** (1995), 207–225.

[Ho] A. Horn, Eigenvalues of sums of Hermitian matrices, *Pacific J. Math.* **12** (1962), 225–241.

[J] S. Johnson, The Schubert calculus and eigenvalue inequalities for sums of Hermitian matrices, Ph.D. thesis, University of California Santa Barbara, 1979.

[Kl] A. A. Klyachko, Stable vector bundles and Hermitian operators, *Selecta Math. (N.S.)* **4** (1998), 419–445.

[Kn] A. Knutson, The symplectic and algebraic geometry of Horn's problem, `math.LA/9911088`, to appear in *Linear Algebra Appl.*

[KT] A. Knutson and T. Tao, The honeycomb model of $GL_n(C)$ tensor products I: Proof of the saturation conjecture, *J. Amer. Math. Soc.* **12** (1999), 1055–1090.

[KTW] A. Knutson, T. Tao, and C. Woodward, The honeycomb model of $GL_n(C)$ tensor products II: Facets of the Littlewood-Richardson cone, in preparation.

[L] B. V. Lidskii, Spectral polyhedron of the sum of two Hermitian matrices, *Funct. Anal. Appl.* **16** (1982), 139–140.

[W] H. Weyl, Das asymptotische Verteilungsgesetz der Eigenwerte lineare partieller Differentialgleichungen, *Math. Ann.* **71** (1912), 441–479.

[Wo] C. Woodward, Honeycomb associativity via scattering theory, preprint.

About the Cover

The cover illustrates one of the principal results from the work of Allen Knutson and Terry Tao. The figure itself is redrawn from one in a preprint of theirs, *The honeycomb model of* $GL_n(\mathbb{C})$. The theorem asserts a relationship between the decomposition of tensor products of representations and a certain collection of what they call "honeycombs". This particular example is concerned with the decomposition of the square of the representation of GL_3 parametrized by the weight vector $(2, 1, 0)$. (The conventions about multiplicities are slightly different in the illustration from what they are in their article—multiplicities of edges are indicated here graphically.) There are some mysteries involved with this and similar pictures—when asked, for example, if there were any direct relationship between the components of such a figure and the decomposition, Knutson and Tao responded, "That's a very good question. We would love to have an interpretation of what the nodes and edge lengths actually mean. For instance, each honeycomb in the picture should correspond to a concrete copy of the appropriate irreducible representation in the tensor product of $V_{2,1,0}$ with itself, with explicit bases and coefficients, etc., but we have no idea how to construct such a canonical decomposition. Nor do we have a particularly good way to enumerate all the honeycombs which are associated with a given tensor product, other than applying off-the-shelf algorithms to enumerate lattice points in polytopes. There is a lot left to be understood in this area."

—Bill Casselman (covers@ams.org)