NOTES FOR THE "GEOMETRIC R-MATRICES" SUMMER SCHOOL

ALLEN KNUTSON

CONTENTS

Part 1. Equivariant Schubert classes

1. EQUIVARIANT K-THEORY OF VARIETIES

Given a commutative ring R, define $K_\bullet$ (aka K-homology, or even "G-theory") using formal differences of finitely generated modules, modulo relations derived from short exact sequences. Then if $R \to S$ and S is R-finite, we get a map $K_\bullet(S) \to K_\bullet(R)$ just by restriction of the action.

Date: December, 2017.

Note that $\otimes_R$ doesn't descend to $K_\bullet$; it doesn't respect SES. So define $K^\bullet$ using only finitely generated projective modules. Now $K^\bullet(R)$ becomes a commutative algebra, acting on $K_\bullet(R)$ as a module. Also, given $R \to S$ we get $K^\bullet(R) \to K_\bullet(S)$, $M \mapsto S \otimes_R M$.

Note that $K_\bullet(R)$ has a "fundamental class" $[R]$ even though $Spec\ R$ is typically noncompact; it will be most analogous to Borel-Moore homology, while $K^\bullet$ is analogous to ordinary cohomology. Acting on this fundamental class, we get a "Poincaré map" $K_\bullet \to K^\bullet$ ("forget that the module's projective"), which is typically neither 1 : 1 nor onto.

Taking Spec, both become functors on affine schemes – now the placement of the $\bullet$ correctly reflects covariance. For general schemes X, instead of modules we use coherent $\mathcal{O}_X$-modules. Pleasantly, we can now define pushforwards in $K_\bullet$ along proper maps, not just finite maps, using sheaf cohomology. For example, under the map $\mathbb{P}^1 \to pt$, we have $[\mathcal{O}(k)] \mapsto k+1 \in K(pt) \cong \mathbb{Z}$.

If G acts on R, even trivially, we can ask for G to act compatibly on modules, and define $K_\bullet^G(R)$, $K^\bullet_G(R)$. First example: if G acts trivially on a point $Spec\ \mathbb{C}$, then $K_G(pt) = Rep(G)$. More specifically, $K_T(pt) \cong Fun(T)$, a Laurent polynomial ring. Note that any X maps to a point, making $K_G(X)$ a $K_G(pt)$-algebra; each pullback $K_G(Y) \to K_G(X)$ is a K_G-algebra morphism; and, each pushforward $K^G(X) \to K^G(Y)$ is a K_G-module morphism. For $\lambda \in T^*$ which we think of additively, let $\exp(\lambda)$ denote the corresponding element of K_T just so that we can think of it multiplicatively.

Example: let $G = T$ act on $X = \mathbb{C}$ with weight $\lambda \in T^*$ (the weight lattice). Then the structure sheaf $\mathcal{O}_{\vec{0}}$ of the origin fits into an equivariant SES

$$0 \to \mathcal{O}_X[-\lambda] \xrightarrow{z\cdot} \mathcal{O}_X \to \mathcal{O}_{\vec{0}} \to 0$$

giving the formula

$$1 - \exp(-\lambda) \in K_T(X) \qquad \mapsto \qquad [\vec{0}] := [\mathcal{O}_{\vec{0}}] \in K_\bullet^T(X)$$

More generally, the class of a subspace $U \leq V$ with weights λ_i on V/U is $\prod_i (1 - \exp(-\lambda_i))$. To "first order in $\{\lambda_i\}$" this is $\prod_i \lambda_i$.

Theorem 1. $K_G(X) = K_T(X)^W$, *a sort of splitting principle, so we usually reduce to* $G = T$.

2. Equivariant localization

Let G be a connected Lie group. Then $K_G(pt) = Rep(G) = Rep(T)^W$ is a domain, so has a fraction field $\mathtt{frac}\ K_G$. Localizing (in the commutative algebra sense) by tensoring with $\mathtt{frac}\ K_G$ is called **equivariant localization**. It fits nicely with the map $K_G(X) \to K_G(X^G) \cong K(X^G) \otimes K_G$:

Theorem 2. *Let* $G = T$*, a torus.*

1. *After equivariant localization, this map is an isomorphism.*
2. *If* X *is smooth projective, this map is injective, even before localizing. Projectivity isn't necessary – it's enough that* T *possess a cocharacter* $\mathbb{G}_m \hookrightarrow T$ *that gives a B-B decomposition (e.g.* X*'s affinization is a cone and* X *is proper over it).*
3. *Still assuming that, the natural map* $K_T(X) \otimes_{K_T} \mathbb{Z} \to K(X)$*, taking each* $\exp(\lambda) \mapsto 1$*, is an isomorphism.*

There's a nice theorem (still assuming, say, smooth projective) that gives the image, but we won't care. (It says: Let X_1 be the union of the 0- and 1-dimensional orbits. Then if a class extends from X^T across the components of X_1, it extends to X.)

One useful formula, that we will abuse later, comes from (1). Assuming X^T finite, every α can be written as $\sum_{f \in X^T} \frac{\alpha|_f}{[f]|_f}[f]$ (proof: their restrictions to each $g \in X^T$ match). Hence the pushforward $p_* \alpha$ of α to a point (which we will just denote $\int \alpha$) is $\sum_{f \in X^T} \frac{\alpha|_f}{[f]|_f}$. This is the Atiyah-Bott '65 "Woods Hole Theorem", which to Fulton's recollection was first proved by Verdier.

We won't take time to deal with equivariant *ordinary* cohomology separately, in part because its definition is much less pleasant than equivariant K-theory's. Also, (2) and (3) only hold if one tensors with $\mathbb{Q}$.

3. Computing with pictures

An extra bonus of being projective, or anyway having a T-equivariant ample line bundle $\mathcal{L}$, is that we can define a "moment map" $\Phi_T : X^T \to T^*, f \mapsto wt(\mathcal{L}|_f)$, and a "moment polytope" $\Phi_T(X) := hull(\Phi_T(X^T)) \subseteq T^* \otimes \mathbb{R}$ (this definition isn't quite right for X not projective – we need to include some directions toward infinity). Also, each component of X^1 has its own moment subpolytope.

Things are especially felicitous if X^T is finite, and easy to draw if Φ_T is injective on X^T (both are true e.g. for ample line bundles on flag manifolds, but only the first is true for $Hilb_n \mathbb{C}^2$). Then we can draw elements of $K_T(X)$ directly on $\Phi_T(X)$, and make use of the dual appearance of T^*.

Example: let $X = \mathbb{P}^2$, with T^2 acting with weights $(0,0), (1,0), (0,1)$; these are also the corners of the moment polytope. The three subvarieties $\mathbb{P}^0, \mathbb{P}^1, \mathbb{P}^2$ (or really their structure sheaves) give a basis of $K_T(\mathbb{P}^2)$. Restricted to fixed points, we get

$$\begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow})(1-e^{\rightarrow}) & - & 0 \end{array} \qquad \begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow}) & - & (1-e^{SE}) \end{array} \qquad \begin{array}{ccc} 1 & & \\ | & \backslash & \\ 1 & - & 1 \end{array}$$

Since it's a basis, we can e.g. compute $[\mathbb{P}^1]^2$. Since restricting to fixed points is a K_T-algebra homomorphism, we can do all computations at the fixed points:

$$\left(\begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow}) & - & (1-e^{SE}) \end{array} \right)^2 = \begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow})^2 & - & (1-e^{SE})^2 \end{array}$$

$$= (1-e^{SE}) \left(\begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow}) & - & (1-e^{SE}) \end{array} \right) + e^{SE} \left(\begin{array}{ccc} 0 & & \\ | & \backslash & \\ (1-e^{\downarrow})(1-e^{\leftarrow}) & - & 0 \end{array} \right)$$

Passing to nonequivariant (setting each $e^\lambda \mapsto 1$), we get $[\mathbb{P}^1]^2 = [\mathbb{P}^0]$ – two lines intersect in a single point.

4. Schubert classes on $P_-\backslash GL_n$

Let $n = \sum_{i=0}^{d} n_i$ be a composition (possibly with 0s), and consider the space of d**-step partial flags** $0 \leq V^{n_1} \leq V^{n_1+n_2} \leq \ldots \leq V^{\sum n_i} = \mathbb{C}^n$, where the latter is *row* vectors and so has a right action of GL_n. The $T = T^n$ =diagonal matrices acts with isolated fixed points, namely where each V^i is coordinate. There are $\binom{n}{n_o,\ldots,n_d}$ of those, which we index by permutations of $0^{n_0}1^{n_1}\cdots d^{n_d}$. **Warning:** this is backwards from most of the literature; we will explain later why we break this convention.

The stabilizer of (each V^k = span of first k coordinates) is block lower triangular matrices P_-. For π a permutation (matrix), let $X_\pi^\circ := P_-\backslash P_-\pi B_+$, a **Bruhat cell** (where B_+ is upper triangular matrices). Then $X_\pi = X_{w\pi}$ for $w \in W_P$, so we should really think of these as indexed by $W_P\backslash W$, or again permutations of $0^{n_0}1^{n_1}\cdots d^{n_d}$. In this case they partition $P_-\backslash G$. The codimension of $X_\pi := \overline{X_\pi^\circ}$ is the number $\ell(\pi)$ of **inversions** in the word π.

Since is T-invariant (right action), it defines an element $[X_\pi]$ of $K_T(P_-\backslash GL_n)$. How can we compute the restriction $[X_\pi]|_\rho$ to a T-fixed point?

If one composition refines another, then $r : P_-\backslash GL_n \twoheadrightarrow Q_-\backslash GL_n$, and by B_+-invariance we know $r(X_\pi) = X_{\pi'}$ for some π', likewise $r^{-1}(X_\rho) = X_{\rho'}$ for some ρ'. Specifically, π' comes from identifying values in π (possibly losing inversions), whereas ρ' comes from breaking ties without introducing inversions. Consequently, it's enough to treat the full-flag case $P_- = B_-$.

Often one wants to think about a Schubert variety X_π geometrically, as those partial flags intersecting the B_+-fixed flag in at least some dimensions. For this it helps to notice that X_π is "boxy"[1] w.r.t. the projections $P_-\backslash GL_n \twoheadrightarrow Gr(k,n)$; it's the intersection of preimages of Schubert varieties in Grassmannians, i.e. the different k-planes are constrained separately. To understand those constraints, coarsen the values in π down to $\{\leq k, > k\}$, and look for descents to see which subspaces in the B_+-flag are relevant.

5. $\check{R}$-matrices appear

We compute $[X_w r_\alpha]$ in the Schubert basis.

First observe that if $w < wr_\alpha$, then $X_w r_\alpha = x_w$.

Otherwise $w > wr_\alpha$. Let $Y := B_-\backslash\overline{B_-wB_+} \times^{B_+} P_\alpha$, where the $\times^{B_+}$ means we divide[2] by the internal diagonal action of B_+. Then we have two P_α-equivariant maps from Y

$$\begin{array}{ccccc} & & Y & \xrightarrow{m} & X_{wr_\alpha} \\ & & \downarrow & & \\ \mathbb{P}^1 & \cong & B_+\backslash P_\alpha & & \end{array}$$

In particular they are $\langle T, r_\alpha\rangle$-equivariant. It is an important fact that $m_*(\mathcal{O}_Y)$ has no higher direct images (the target has "rational singularities"), so $m_*[Y] = [m(Y)] = [X_{wr_\alpha}]$.

[1] One of the big theorems in flag matroids is that a subset $M \subseteq S_n$ with the Coxeter matroid property "each $\pi \cdot M$ has a unique Bruhat minimum" is similarly boxy for the matroid projections $S_n \twoheadrightarrow \binom{n}{k}$.

[2] Fiber products, which are subsets, are denoted by $\times_S$; this is rather a quotient so we use $\times^S$.

On this $\mathbb{P}^1$ we have a single relation between $[\mathbb{P}^1]$, $[0]$, $[\infty]$ which we can compute from their restrictions to fixed points:

	$[\mathbb{P}^1]$	$[0]$	$[\infty]$
restriction to 0	1	$1-\exp(\alpha)$	0
restriction to ∞	1	0	$1-\exp(-\alpha)$

$$\therefore \qquad [\infty] = (1-\exp(-\alpha))[\mathbb{P}^1] + \exp(-\alpha)[0]$$

Sanity check: nonequivariantly this become $[\infty] = [0]$.

Pull back this equation along the vertical map to Y:

$$[X_w]r_\alpha = (1-\exp(-\alpha))[Y] + \exp(-\alpha)[X_w] \qquad \in K_T(Y)$$

Push forward along m:

$$[X_w]r_\alpha = (1-\exp(-\alpha))[X_{wr_\alpha}] + \exp(-\alpha)[X_w]$$

Let's write this down for the 2^2 Schubert classes on $Gr(\mathbb{C}^2) = Gr_0(\mathbb{C}^2) \coprod Gr_1(\mathbb{C}^2) \coprod Gr_2\mathbb{C}^2)$, where $\alpha = y_1 - y_2$ in T^2-coordinates:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ \\ 0 & \exp(y_1-y_2) & 1-\exp(y_1-y_2) & 0 \\ 0 & 0 & 1 & 0 \\ \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

a 5-vertex model. R-matrix people will tell you that you should seek instead a 6-vertex model, and we'll hear more about that.

6. Subword complexes and the AJS/Billey and Graham/Willems formulæ

Let Q be a word in S_n's generators, and $\Delta^{|Q|-1}$ the simplex whose vertices are indexed by the letters in Q – or better yet, the subwords of Q *missing* a single letter. Then the other faces also correspond to subwords, by intersecting. Define $m : 2^Q \to W$ by $F \subseteq Q$ maps to the **Demazure product** of $Q \setminus F$, the unique Bruhat-largest $\prod G$, $G \cap F = \emptyset$. Then

Theorem 3. (1) *Each $m^{-1}(w)$ is homeomorphic to an open ball (or very rarely a sphere), as a locally closed subcomplex of $\Delta^{|Q|-1}$.*
(2) *Its closure $\Delta(Q, w)$, the* **subword complex**, *is homeomorphic to a ball or sphere.*
(3) $\partial\Delta(Q, w) = \bigcup_{w'>w} \Delta(Q, w')$.
(4) *Consequently, the Möbius function on $\Delta(Q, w)$ is ± 1 on those faces F that enjoy $m(Q \setminus F) = w$ on the nose. (In general simplicial complexes, $\mu(F) = 1 - \chi(\mathrm{link}(F))$.)*

I like to express this as m defines a "Bruhat decomposition of $\Delta^{|Q|-1}$".

Stanley-Reisner theory is largely the observation that simplicial complexes are equivalent data to unions of coordinate subspaces in $\mathbb{C}^Q$. If we let T act on $\mathbb{C}^Q$ by with weight $\prod_{j<i} r_{q_j} \cdot \alpha_{q_i}$ on the ith coordinate,[3] then we can state the Graham/Willems formula for

[3] These roots are all positive roots if Q is a reduced word.

the restriction:

$$[X_\pi]|_\rho \leftrightarrow \sum_{F\in\Delta(Q,\pi)^\circ} (-1)^{\mathrm{codim}}[\mathbb{C}^F] \qquad \in K_T(\mathbb{C}^Q) \cong K_T$$

where Q is any reduced word for ρ. It's easy to derive from the $\check{R}$-matrix result from the last section.

In view of (4), though, I observed that the RHS is just $[SR(\Delta(Q,\pi))]$.

It may seem sad to have to choose a word Q for σ given that the answer doesn't depend on it. Certainly commuting moves on Q do not really change the formula, only braid moves. In lucky "fully commutative" cases of σ, such as σ 321-avoiding (in type A), only commuting moves are possible and hence the formula is essentially unique. This is nearly the context of the Ikeda-Naruse "excited diagrams" formula.

The Graham/Willems formula simplifies, to an older formula of AJS/Billey, if we limit from K_T to H_T. Then (1) only the facets (maximal faces) of the subword complex contribute, and (2) the individual terms $\prod(1-e^{-\beta}) \in K_T$ limit to $\prod \beta \in H_T$.

7. Gröbner geometry of the Graham/Willems formula

Towards explaining the $[SR(\Delta(Q,\pi))]$ simplification, define $X^\rho_\circ := B_-\backslash B_-\rho B_-$, and factor $\{\rho\} \hookrightarrow B_-\backslash GL_n$ as

$$\begin{array}{ccccc} & & X^\rho_\circ \cap X_\pi & \hookrightarrow & X_\pi \\ & & \downarrow & & \downarrow \\ \{\rho\} & \longrightarrow & X^\rho_\circ & \longrightarrow & B_-\backslash GL_n \\ \uparrow & & \uparrow m & & \\ \vec{0} & \longrightarrow & \mathbb{A}^{\ell(\rho)} & & \end{array}$$

This helps because the intersection is transverse, so the K-cohomological pullback of $[X_\pi] \in K_T(B_-\backslash GL_n) \to K(X^\rho_\circ)$ is just the class $[X_\pi \cap X^\sigma_\circ]$. Then the second pullback $K_T(X^\sigma_\circ) \to K_T(\{\mu\})$ is trivial, as both K_T-algebras are just K_T.

The reduced word for ρ enters in actually putting coordinates on $X^\rho_\circ$, via the **open Bott-Samelson map**

$$m : (z_1,\ldots,z_{|Q|}) \mapsto P_-\backslash P_- \prod_{i=1}^{|Q|} \begin{bmatrix} 1 & & & & & & \\ & \ddots & & & & & \\ & & 1 & & & & \\ & & & z_i & 1 & & \\ & & & -1 & 0 & & \\ & & & & & 1 & \\ & & & & & & \ddots \\ & & & & & & & 1 \end{bmatrix} \text{in position } q_i \in [1,n)$$

Theorem 4. *Replace each of the equations defining $m^{-1}(X_\pi)$ by their lexicographic initial terms (this "Gröbner degeneration" doesn't change the K_T-class). Then the result is the Stanley-Reisner ideal of $\Delta(Q,w)$. This degeneration commutes with taking unions and intersections of $\{X_\pi\}$, as well.*

8. A GLOBAL VIEW IN TYPE A: GROTHENDIECK POLYNOMIALS

Let $P_- \leq GL(2n)$ be $n+n$ block lower triangulars. Fulton '92 considers the maps

$$M_{n\times n} \xrightarrow{\text{graph}} Gr(n,2n) \leftarrow P_-\backslash GL(2n) \leftarrow X_\circ^{w_0 w_0^P}$$

and observes that the resulting isomorphism, when fit into

$$B_-\backslash GL_n \leftarrow GL_n \hookrightarrow M_n \cong X_\circ^{w_0 w_0^P},$$

lets one see the Schubert strata X_π, $\pi \in S_n$ of $B_-\backslash GL_n$ using the strata $X_{\pi\oplus 1_n} \cap X_\circ^{w_0 w_0^P}$.

Specifically, define the **double Grothendieck (Laurent) polynomial**

$$G_\pi(x,y) := [X_{\pi\oplus 1_n}]|_{w_0 w_0^P}$$

which one can compute from the Graham-Willems formula. The choice of reduced word Q for $w_0 w_0^P$ is essentially unique ($w_0 w_0^P$ is "fully commutative"), and the Graham/Willem summands correspond to "nonreduced pipe dreams" for π.

One can compute point restrictions using these: $[X_\pi]|_\rho = G_\pi(\rho\cdot y, y)$. To see this, first restrict the $T^n \times T^n$ acting on $M_{n\times n}$ to the ρ-twisted diagonal (accounting for the specialization), then restrict to $\overline{T\sigma}$, and follow under the isomorphisms $K_T(B_-\backslash GL_n) = K_T(T\backslash GL_n) = K_{T\times T}(GL_n)$.

Part 2. Schubert calculus

Now that we have the classes $\{S_\lambda := [X_\lambda]\}$, in principle we can compute the multiplication of two Schubert classes and re-expand in the basis, with coefficients from K_T.

$$S_\lambda S_\mu = \sum_\nu c_{\lambda\mu}^\nu S_\nu$$

Since $S_\lambda|_\nu = 0$ unless $\nu \geq \lambda$, we can straightaway learn that $c_{\lambda\mu}^\nu = 0$ unless $\nu \geq \lambda, \mu$. Remember for later the ν = identity case: $c_{\lambda\mu}^{Id} \neq 0 \implies \lambda = \mu = Id$ and $c_{\lambda\mu}^{Id} = 1$.

In particular, it is very easy to find $c_{\lambda\mu}^\nu$ by computer for small examples. If we want to compute in nonequivariant K-theory, set all $y_i = 1$; if we want to compute in ordinary equivariant cohomology, expand to first order in y_i and only consider terms of the correct y-degree $\ell(\nu)-\ell(\lambda)-\ell(\mu)$. (One can do this with the classes first, and only go to structure constants afterward.)

$$n = 1: \qquad S_0^2 = S_0 \qquad S_1^2 = S_1$$

$$n = 2: S_{00}^2 = S_{00} \qquad S_{01}^2 = S_{01} \quad S_{01}S_{10} = S_{10} \quad S_{10}^2 = (1-y_2/y_1)S_{10} \qquad S_{11}^2 = S_{11}$$

In general if a word λ with content $0^{n_0}1^{n_1}\cdots d^{n_d}$ is the weakly increasing one, then $S_\lambda|_\mu = 1\ \forall\mu$, so $S_\lambda = [P_-\backslash GL_n]$ is the identity for multiplication.

9. THE ORACLE SPEAKS

A **puzzle piece** will be a size 1 equilateral triangle (oriented either Δ or ∇) with edges[4] labeled, from some fixed label set. A **size** n **puzzle** will, at first, be a size n triangle built from n^2 pieces ($\binom{n+1}{2}$ Δs and $\binom{n}{2}$ ∇s) glued so that edge labels agree. Later another shape will be added.

> **Oracle:** there is an edge label set $[0, d]$ plus others, and a finite set of allowed puzzle pieces, such that nonequivariant ordinary cohomology $c_{\lambda\mu}^{\nu}$ is the number of size n puzzles with labels λ, μ, ν on the NW, NE, S sides (read left to right). In particular, the extra labels only appear in the interior.

Why should there be this Z_3 rotational symmetry?

$$c_{\lambda\mu}^{\nu} = \int_{P_-\backslash GL_n} S_\lambda S_\mu S^\nu = \int_{P_-\backslash GL_n} S_\lambda S_\mu S_{\nu \text{ backwards}}$$

The fact that the dual basis $\{S^\nu\}$ (with respect to the multiply-then-integrate pairing) is a reordering of the original basis is special to ordinary cohomology, not K- or equivariant.

In fact the oracle only told me this[5] for $d = 1$, and I incorrectly generalized to all d; it turns out to only hold like this for $d = 1, 2, 3$.

From $S_0^2 = S_0$ and $S_1^2 = S_1$, we learn that there must be a $(0, 0, 0)$-triangle and $(1, 1, 1)$-triangle. From $S_{\overline{01}} S_{01}$ we want a puzzle with labels

$$\begin{array}{c} 1/\backslash 0 \\ 0/--\backslash 1 \\ 01 \end{array}$$

We know pieces that fit in the SW and SE corners, so let's try putting them in; what remains is

$$\begin{array}{c} 1/\backslash 0 \\ 0\backslash/1 \end{array}$$

so if we invent a new label "(10)" to go on that middle edge, and create this new piece (and its rotation), then we've accounted for that coefficient.

Theorem 5 (KTW '03)**.** *The oracle was right, with these three labels* 0, 1, (10) *and these three puzzle pieces (up to rotation).*

Our proof went via bijection to a known rule (Berenstein-Zelevinsky triangles).

The $d = 2$ case is easy to explore experimentally. Since it includes the $d = 1$ case we'll need the labels $0, 1, 2, (21), (20), (10)$, and these suffice for $n \leq 2$ (where we can only fit $d = 1$ anyway). At $n = 3$ we're forced to have two new labels $2(10)$, $(21)0$ still on triangles $Y, X, (YX)$ clockwise, for a total of 8. (For example, $S_{021} S_{102}$ is about flags $L^1 < P^2$ such that $P \geq \mathbb{C}^1$ and $L \leq \mathbb{C}^2$. Hence either $L = \mathbb{C}^1$, or $P \geq L + \mathbb{C}^1 = \mathbb{C}^2$, giving $S_{201} + S_{120}$; check for yourself that these two puzzles suggest one introduce $(21)0$, $2(10)$ respectively.)

I conjectured in '99 that those 8 labels and pieces compute nonequivariant ordinary Schubert calculus on 2-step flag manifolds; this was proven in [BKPT14].

[4]It is tempting, when drawing just one of them, to put the edge labels outside the triangle; this should be avoided since it becomes impossible when the pieces are glued together.

[5]Terry Tao and I invented puzzles to solve Weyl's 1912 Hermitian sum problem. Klyachko solved it with Grassmannian Schubert calculus. So we wondered, are they directly related?

Again, since 3-step includes 2-step, we'll need $(ji), k(ji), (kj)i$ labels for $k > j > i$. I conjectured that any binary tree with strictly decreasing leaves (like $(32)(10)$) should give a label and piece, but these 23 such define a noncommutative multiplication. Buch found the four missing labels, e.g. $3(((32)1)0)$, in which a weak decrease is allowed if separated by three parentheses (!?!).

Paul Zinn-Justin and I proved this 3-step, 27 label rule this year, using E_6's R-matrix acting on $27 \otimes 27$.

10. Ř-MATRICES REAPPEAR

The equivariant ordinary cohomology $c^\nu_{\lambda\mu}$ have to be polynomials in $y_i - y_j$, and again it is easy to compute them for small n via equivariant localization:

$$S_{10}^2 = (y_2 - y_1) S_{10}$$

Note that there is no longer any Z_3-symmetry[6] , so we shouldn't need to insist that everything be triangles. Since our answer has $y_i - y_j$ factors in it, we aren't just counting any more; some pieces may be worth a factor, depending on their location.

For $d = 1$, it turns out (KT '03) to be enough to invent an "equivariant rhombus" $\begin{smallmatrix} 0/\backslash 1 \\ 1\backslash/0 \end{smallmatrix}$ worth $y_i - y_j$ if it's in the ith NW/SE diagonal and jth NE/SW diagonal. For $d = 2$, there are a couple more equivariant rhombi [Buch 2015]. For $d = 3$, subject to some assumptions, there is no equivariant extension of the 27-label rule from before.

Because of the rhombi, it's no longer natural to think of the puzzle as made of n^2 triangles, but rather $\binom{n}{2}$ vertical rhombi and n Δs along the bottom. It will be more convenient to work with the **dual graph picture**, a superposition of n Y-vertices laid left to right. We think of every edge as oriented downward, and the n Ys as colored $y_1, \ldots, y_n$ from left to right. In fact we'll actually color them too: green from NW, red from NE, blue South.

If we think of the edge labels (say $0, 1, (10)$) as indexing bases of green and red vector spaces V_g, V_r (over $\mathrm{frac}\ H_T$), then we can think of a tetravalent vertex as giving a map $V_g \otimes V_r \to V_r \otimes V_g$. In fact for $d = 1, 3, 4$ we'll have $V_g \cong V_r$.

In ZJ'06 Paul observed that these $d = 1$ tetravalent vertices satisfy the Yang-Baxter equation, i.e. this matrix $V_g \otimes V_r \to V_r \otimes V_g$ is an Ř-matrix. I don't know any reasonable way to retrodict this, since (recast in terms of rhombi) it involves equivariant rhombi placed in two illegal directions.

11. MIXING THE TWO Ř-MATRICES

To prove that puzzles are indeed computing $(c^\nu_{\lambda\mu})$, i.e. that $S_\lambda S_\mu = \sum_\nu c^\nu_{\lambda\mu} S_\nu$, it suffices to prove

$$S_\lambda|_\sigma S_\mu|_\sigma = \sum_\nu c^\nu_{\lambda\mu} S_\nu|_\sigma$$

for each T-fixed point σ. Recall that we have a nice formula for $S_\nu|_\sigma$ using a reduced word for σ, which we can now interpret in the dual graph picture as a vertical wiring

[6]Very strangely, nonequivariant K-theory of Grassmannians or other minuscule (not cominuscule) flag manifolds *does* enjoy this Z_3-symmetry, because $X^\nu = X_{\nu \text{ backwards}}(1 - X_\square)$, so $c^\nu_{\lambda\mu} = \int S_\lambda S_\mu S^\nu = \int S_\lambda S_\mu S_{\nu \text{ backwards}}(1 - S_\square)$.

diagram with ν across the top and the identity $\mathtt{sort}(\nu)$ across the bottom. In particular, the $\check{R}$-matrix from §10 reappears as an endomorphism of $V_b \otimes V_b$.

This has the intriguing benefit that the RHS can now be interpreted as the amplitude of a *single* graph (with boundary labeled λ, μ,identity), no external sum over ν necessary. (A priori we might worry that the ν edges in the graph might admit labels other than $[0, d]$; one must prove separately that this doesn't happen.)

Not only $\check{R}$-matrices are here; the trivalent vertices are maps $V_g \otimes V_r \to V_b$. So the entire diagram gives a map $V_g^{\otimes n} \otimes V_r^{\otimes n} \to V_b^{\otimes n}$, and λ, μ, ν indicate which matrix entry to look at.

12. Ingredients of the proof

The YBE that Paul checked lets us manipulate diagrams like these; assume that all such isotopies give similar tensor identities. (These are identities between $27^3 \times 27^3$ matrices of rational functions, so this is not a trivial step! We will come back to it in the next section.)

Then we can move the blue/blue crossings in σ's reduced word up through the puzzle region, to become red/red crossings *and* green/green crossings (they're not conserved).

Now the puzzle region has the identity across the South side. If one *believes* the puzzle rule, then one expects that the only filling of the puzzle region is identity * identity = $1 \cdot$ identity. But at this point in the proof, one must give a separate argument for this.

Now the puzzle region can be completely ripped out of the diagram, leaving a green/green σ-crossing diagram going from λ to the identity, similarly a red/red, giving exactly the LHS of the identity we wanted.

13. An origin of the YBE and bootstrap equations

The quantized loop algebra $U_q(\mathfrak{g}[z, z^{-1}])$ has many finite-dimensional representations $V_{\lambda,c}$ indexed by a dominant weight λ and a complex parameter c (and many others too!). If λ is minuscule, then as reps of the subalgebra $U_q\mathfrak{g}$, these are just the irrep V_λ.

For generic values c, c', the tensor product $V_{\lambda,c} \otimes V_{\mu,c'}$ is again irreducible, and nonobviously isomorphic to $V_{\mu,c'} \otimes V_{\lambda,c}$; the Schur's lemma isomorphism depends on c/c' and gives a trigonometric $\check{R}$-matrix.

However, we need the trivalent vertex $V_g \otimes V_r \to V_b$, so the parameters must *not* be generic. Specifically, we need the n parameters y'_i on the red edges to be q^2 times the n parameters y_i on the green edges.

$d = 1$. $\mathfrak{g} = \mathfrak{sl}_3 = A_2$, $V_g = V_r = \mathbb{C}^3$, $V_b = \mathrm{Alt}^2\mathbb{C}^3$. The $\mathfrak{g}$-decomposition of $V_g \otimes V_r$ is $\mathrm{Alt}^2\mathbb{C}^3 \oplus \mathrm{Sym}^2\mathbb{C}^3$. At the special value of c/c', the $\mathrm{Sym}^2\mathbb{C}^3$ becomes a subrepresentation and $\mathrm{Alt}^2\mathbb{C}^3$ becomes a quotient. (At the reciprocal value the reverse is true.)

Note that if one pursues this, one ends up with the 6-vertex model, not the 5-vertex we needed, so we still need to do some degeneration.

$d = 2$. $\mathfrak{g} = \mathfrak{spin}_8 = D_4$, $V_g = \mathrm{spin}_+$, $V_r = \mathrm{spin}_-$, $V_b = \mathbb{C}^8$. Here the Z_3-invariance involves the triality of D_4.

$d = 3$. $\mathfrak{g} = E_6$, $V_g = V_r = 27$, $V_b = 27^*$.

The degeneration involved essentially multiplies some columns of the $\check{R}$-matrix by powers of q and rows by the opposite powers, then takes $q \to 0$. For $d = 1, 2$ everything works beautifully and we recover puzzle rules in K_T. For $d = 3$ some terms go to infinity, and we can only get a limit if we first give up equivariance, resulting in a K-rule (hence H-rule) but no K_T- or H_T-rules.

$d = 4$. Here we know the correct analogues are $\mathfrak{g} = E_8$, $V_g = V_r = V_b = \mathfrak{e}_8 \oplus \mathbb{C}$. Much goes wrong in this nonminuscule situation.

Part 3. A_d **quiver varieties**

If we don't want to take $q \to 0$ (as we sort of don't for $d = 3$ and really don't for $d = 4$) then we have to move beyond d-step flag manifolds to Nakajima quiver varieties.

14. Definition and (our) main example

Take a simply-laced Dynkin diagram (soon to be A_d, d vertices in a linear graph), attach an extra "framed" vertex to each of the old "gauged" vertices, and put arrows both ways along each edge; this is a **Nakajima quiver**.

Label the framed vertices with finite-dim vector spaces, and consider that dimension vector as giving a linear combination λ of A_d fundamental weights $\{(1, \ldots, 1, 0, \ldots, 0)\}$ (length $d + 1$, and not all 0 or 1). Equivalently, this is a choice of dominant weight.

Then label the gauged vertices, but consider this as determining $\mu = \lambda$ minus a linear combination of simple roots $\{(0, \ldots, 0, 1, -1, 0, \ldots, 0)\}$. In particular, if μ is a weight in the irrep V_λ, then it is of this form.

Now we want the quiver representations (choice of linear map, for each arrow) satisfying the "moment map condition" that at each *gauged* vertex, $\sum \pm$"go out, then come back in"$= 0$. (The largely unimportant signs are determined by an orientation of the singly-oriented Nakajima quiver.) When $\lambda = 0$ these are called "preprojective" conditions.

After having imposed those conditions, pass to a certain "stable" open set I won't describe, then divide by $\prod GL$(the gauged vector spaces), to obtain the smooth irreducible holomorphic symplectic **quiver variety** $\mathcal{M}(\lambda, \mu)$. Let $\mathcal{M}(\lambda) := \coprod_\mu \mathcal{M}(\lambda, \mu)$ which I will call a **quiver scheme** just because it's reducible.

To name points in such a quotient, one has to work in the framed vertices, looking at paths from one framed vertex to another, images of paths from gauged vertices to framed, or kernels of paths from framed vertices to gauged.

Our main example: $\lambda = n\omega_1$ for A_d. Let $E \in End(\mathbb{C}^n)$ be the composite $\mathbb{C}^n \to \mathbb{C}^{\mu_1} \to \mathbb{C}^n$.

$$\begin{aligned} E^2 &= (\mathbb{C}^n \to \mathbb{C}^{\mu_1} \to \mathbb{C}^n \to \mathbb{C}^{\mu_1} \to \mathbb{C}^n) \\ &= -(\mathbb{C}^n \to \mathbb{C}^{\mu_1} \to \mathbb{C}^{\mu_2} \to \mathbb{C}^{\mu_1} \to \mathbb{C}^n) \end{aligned}$$

Continuing this way,

$$E^d = \pm(\mathbb{C}^n \to \cdots \to \mathbb{C}^{\mu_d} \to \cdots \to \mathbb{C}^n)$$

and $E^{d+1} = 0$. Moreover, if we define $V_i = image(\mathbb{C}^{\mu_{d+1-i}} \to \mathbb{C}^n)$, then the (V_i) are a partial flag satisfying $EV_i \leq V_{i-1}$. The undescribed stability condition turns out to be that $V^{\mu_i} \to$

$\mathbb{C}^n$ is injective, so the data $(E, V^\bullet)$ is an element of a well-defined Springer resolution T^*Fl. Finally, one need argue that this map $\mathcal{M}(n\omega_1, \mu) \to T^*Fl$ is an isomorphism.

15. Applications of quiver varieties

Why make these?

Theorem 6. (1) *(Kronheimer)* $\mathcal{M}(\textit{highest root}, 0) \cong \widetilde{\mathbb{C}^2/\Gamma}$.
(2) *(Nakajima)* $H_{top}(\mathcal{M}(\lambda))$ *supports a* $U\mathfrak{g}$*-rep (defined using convolutions) making it* V_λ.
(3) *(Varagnolo)* $H_*(\mathcal{M}(\lambda))$ *supports a* $U_q(\mathfrak{g}[z])$*-rep.*
(4) *(Nakajima)* $K(\mathcal{M}(\lambda))$ *supports a* $U_q(\mathfrak{g}[z^\pm])$*-rep.*
(5) *(Yaping et al., unifying those two)* $E(\mathcal{M}(\lambda))$ *supports a* $U_q(\mathfrak{g} \otimes E_{\mathbb{C}^\times}(pt))$*-rep.*

Recall that Jimbo et al. made R-matrices using $U_q(\mathfrak{g}[z^\pm])$-reps. Taking the pushout of that with Nakajima's second result above, we expect to see R-matrices in K(quiver schemes), which Andrei will talk more about.

Consider two dimension vectors λ and λ', and the "direct sum" map

$$\oplus : \mathcal{M}(\lambda) \times \mathcal{M}(\lambda') \to \mathcal{M}(\lambda + \lambda').$$

"Being a direct sum" is the same data as "having a circle action with weights 0 and 1", so the target quiver variety has such an action, and indeed the image of the above map is the inclusion of fixed points. It turns out to be reasonable to expect $H(\mathcal{M}(\lambda) \times \mathcal{M}(\lambda'))$ and $H(\mathcal{M}(\lambda + \lambda'))$ to be isomorphic, and part of Andrei and Davesh Maulik's stable envelope construction is to set up an $U_q(\mathfrak{g}[z])$-equivariant isomorphism.

Consider the example of the quiver scheme $\mathcal{M}(n\omega_1)$. Each component is a cotangent bundle, hence its H_{top} is one-dimensional, as befits the the weight spaces of the irrep $V_{n\omega_1} = Sym^n(\mathbb{C}^{d+1})$. Whereas the total homology is $(\mathbb{C}^{d+1})^{\otimes n}$; regarding it as an n-fold tensor product corresponds to breaking λ as a sum of n terms (each ω_1), allowing us to compute T^n-equivariant homology of $\mathcal{M}(n\omega_1)$. The resulting equivariant parameters then serve as the (generic) parameters of the $U_q(\mathfrak{g}[z])$-irrep tensor factors. Finally note that this $(d+1)^n$ breaks as a sum of multinomial coefficients, which are the dimensions of homology of the individual quiver varieties.

16. Deformations of quiver varieties, and a first approach to Maulik-Okounkov stable classes

There are three[7] avenues of deformation, which we spell out in the $\mathcal{M}(n\omega_1, \mu)$ case.

(1) Instead of imposing $\Phi_{\prod GL} = 0$ (the moment map condition), we could impose that $\Phi_{\prod GL} = (\tau_i \mathbf{1})$, an independent scalar at each vertex. This corresponds to taking a Springer resolution of another principal orbit closure (one Jordan block for each eigenvalue).
(2) We could change the (still undescribed) stability condition. For example, instead of taking the Springer resolution, we could just get the (singular) nilpotent cone.
(3) We could deform in a noncommutative direction, which in the case of a cotangent bundle would move to differential operators on the base.

[7] At least three. There also exist "multiplicative quiver varieties", for example, which we won't touch here.

In particular, consider the following sheaves of (noncommutative) algebras over G/B_+. Let μ be a general element of $\mathfrak{t}^*$, thought of (using the Killing form) as a subset of $\mathfrak{g}^*$.

$$\begin{array}{ccc} \mathcal{D}_\mu & \dashrightarrow & \mathcal{D} = \mathcal{D}_0 \\ \downarrow & \text{degenerations, all} & \downarrow \\ \mathcal{O}_{G\cdot\mu \cong G/T} & \dashrightarrow & \mathcal{O}_{T^*G/B} \end{array}$$

The most general is $\mathcal{D}_\mu$, μ-twisted differential operators, whose modules (under Beilinson-Bernstein equivalence) match with category $\mathcal{O}$ at central character given by μ. That has $|W|$ Verma modules, which are irreducible and have no Exts. Defining "Verma" already involves making a choice of Borel, and relating two choices involves R-matrices.

When we go commutative (down), we get the smooth space G/T, with $|W|$ T-fixed points. The choice of Borel determines a dominant cocharacter of T, with which to define a Białynicki-Birula decomposition of... only part of G/T. The BB strata, $\{BwT/T\}$, all are closed and Lagrangian (and don't come close to covering G/T).

If we instead went to $\mu = 0$, we'd get category $\mathcal{O}$ at dominant integral (zero) weight, where the Vermas have a complicated Kazhdan-Lusztig relation to the irreducibles, plus lots of Exts.

Now finally we go to T^*G/B, where the Lagrangians $\{BwT/T\}$ have become cotangent-dilation-invariant, so defining classes in $H^*_{T\times\mathbb{C}^\times}(T^*G/B) \cong H^*_{T\times\mathbb{C}^\times}(G/B) \cong H^*_T(G/B)[\hbar]$. These are the **Maulik-Okounkov[8] "stable" classes.** They give a basis for $H^*_G(G/B)[\hbar]$, once one inverts $\hbar$.

One hint that these are better than usual Schubert classes is that they are related by the operators $r_\alpha + \hbar\partial_\alpha$, which square to 1 instead of to 0; eventually this will say that if we did the calculation from §10 with these, we'd get the (rational) 6-vertex model not 5-vertex.

Changjian Su gave a formula for the point restrictions of M-O stable classes, which implies the following one with complicated multiplicities:

$$[MO_\lambda]|_\mu = (\text{factor}) \sum_{F\in\Delta(Q,\lambda)} \#\left\{ I \supseteq F \,:\, \prod(Q\setminus I) = w \right\} \; [\text{conormal variety to } \mathbb{A}^F]$$

(Even though this is in H_T, it's a sum over all faces of the subword complex, not just facets.) Su's version is a clever regrouping of these terms resulting in no multiplicities, but also no clear geometry. It would be very interesting to find an analogue of §7 explaining this formula.

17. Retrodicting the $d = 1$ puzzle rule

In linear algebra terms, we started with a basis vector in $V_g^{\otimes n} \otimes V_r^{\otimes n}$, do a bunch of $\check{R}$-matrices to it until it was in $(V_g \otimes V_r)^{\otimes n}$, then did n trivalent vertices to that to end up in $V_b^{\otimes n}$.

Both V_g, V_r were the (top=total) homology of the A_2 Nakajima quiver scheme $\mathcal{M}(\omega_1)$. However, the $0, 1, 10$ basis vectors in V_g are Z_3-permuted from the $0, 1, 10$ basis vectors in V_r (this is an effect of our wanting the trivalent tensor to be $\mathfrak{g}$-invariant, and in particular

[8] In the $\mathcal{M}(n\omega_1)$ situation, these match the much older "Chern-Schwartz-MacPherson classes of the Bruhat cells".

for the $(0,0,0)$-triangle term to have vanishing t-weight). Specifically, the $0^k 1^{n-k}$ on the NW and NE sides give us basis vectors with different weights, from the quiver varieties

$$\mathcal{M}\begin{pmatrix} n & & \\ | & & \\ k & - & 0 \end{pmatrix} \quad \text{and} \quad \mathcal{M}\begin{pmatrix} n & & \\ | & & \\ n & - & k \end{pmatrix}.$$

(Note that these are A_1 quiver varieties disguised as A_2, which has to do with our only using $0, 1$ on the boundary not (10).)

Now we combine these to get a $\mathcal{M}\begin{pmatrix} 2n & & \\ | & & \\ n+k & - & k \end{pmatrix}$ quiver variety, which is an honest 2-step flag manifold's cotangent bundle. That is, we have a Maulik-Okounkov stable Lagrangian therein.

When we've applied some of the R-matrices/tetravalent vertices, it's still about the same cotangent bundle – just about changing the cocharacter used to define the BB strata.

Then finally we have a stable class on this cotangent bundle of $(0 \leq V^k \leq V^{n+k} \leq \mathbb{C}^{2n})$ flags, of dimension $2(k(2n-k)+n(n-k))$, but we want one in $T^*Gr(k,n)$ of dimension $2k(n-k)$; the difference is $2n^2$. Note too that the first space has a (framed vertex) action of $GL(2n)$, the second only $GL(n)$, and finally that we expect to break the T^{2n} inside that $GL(2n)$ down to something n-dimensional (since the spectral parameters must be related). We believe that the right move is to take the symplectic quotient by $Rad(P_{n+n})$ (work in progress).

E-mail address: `allenk@math.cornell.edu`