NOTES FOR MINICOURSE ON SCATTERING AMPLITUDES AND REPRESENTATION THEORY

1. The generic decomposition fan

Today, rather than getting right into the details of my topic, I want to motivate and introduce somewhat informally the g-fan, which will also play an important role in other mini-courses this week as well.

1.1. A_2. Let's start with a very simple quiver $1 \xleftarrow{\alpha} 2$. A representation of this quiver is just a pair of vector spaces X_1, X_2, and a linear map between them. The *dimension vector* of the representation just records $(\dim X_1, \dim X_2)$.

We want to ask a rather simple-seeming question. What does a typical representation of this quiver with dimension vector (d_1, d_2) look like? (Of course, we will have to sort out what "typical" means.)

Let's start simple. If the dimension vector is $(1, 0)$, the linear map must be the zero map, and so the representation is just the simple representation at 1, which we denote S_1, with no choices at all.

How about the dimension vector $(1, 1)$? Well, here, the representation looks like

$$\mathbb{C} \xleftarrow{X_\alpha} \mathbb{C}$$

If X_α is zero, we get a representation isomorphic to the direct sum of the simple modules we already saw at vertices 1 and 2.

For all values of α non-zero, the resulting representations are isomorphic, since they differ by a change of basis. This is a third indecomposable representation of this quiver. And this is the one we'll consider the "typical" representation of this dimension vector. (I think we can agree that "non-zero" is more "typical" than "zero".) Let's call it $M_{1,1}$.

Now let's consider what happens at dimension (d, d), for d positive. The representation is

$$\mathbb{C}^d \xleftarrow{X_\alpha} \mathbb{C}^d$$

and X_α is a $d \times d$ matrix. What does a typical $d \times d$ matrix look like?

It's time to be a little more precise about what we mean by a "typical" representation. To specify a representation, we have to specify the entries of (usually) several matrices (just one in this case). The space of choices is just a big affine space, which we denote $\mathrm{Rep}_Q((d, d))$. (In this case, of dimension d^2.) So we can define a "typical" behaviour to be a behaviour that holds on the complement of a region cut out by a finite number of algebraic equations. (We are looking for a "Zariski open" subset of the affine space where the representations are all isomorphic, if that term is familiar.) We call this the *generic* behaviour. Nothing, by the way, so far guarantees that there will be an isomorphism class which has this property of holding generically. Nonetheless, that is what we are hoping for, and which will turn out to happen in the cases we are looking at today.

When the determinant of this $d \times d$ matrix is non-zero (which is an condition of the kind we want: complement of an algebraic equation) the matrix will be invertible. That means that, after a change of basis at source and target, it can be changed into an identity matrix, and the representation will be isomorphic to $M_{1,1}^d$.

Let's now consider the case where our dimension vector is (d_1, d_2) with $d_2 < d_1$. What does a typical $d_2 \times d_1$ matrix look like? Well, generically (outside the set given by the vanishing of another determinant), the first d_2 rows form an invertible matrix, which implies that the matrix as a whole defines an injective map. As a result, we're going to get d_2 copies of $M_{1,1}$ and $(d_1 - d_2)$ copies of S_1. (In fact, the map can still be injective even if this determinant is zero. The details are left for the exercises.) The results for $d_1 > d_2$ are similar.

How should we summarize our results? Let's look at the classes of dimension vectors which all have the same summands appearing in their generic decomposition.

There are six different posssibilities which we can summarize in the following picture:

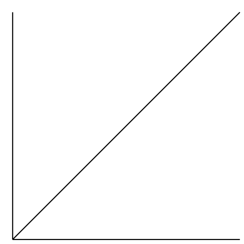

If the dimension vector is on one of the rays, then its generic decomposition is just a sum of copies of the representation whose dimension lies on the ray, while if a dimension vector lies between two rays, then its generic decomposition is sums of the representations corresponding to the two rays. (And if the dimension vector is zero, its generic decomposition is zero.)

I'm going to call this collection of cones the "generic representation fan."

1.2. **Dynkin type.** The amazing thing is that this generalizes beautifully to all Dynkin types (and even beyond them). An excellent source for this material, which goes into a lot more detail than I can in a single hour, are the notes by Brion [**?**].

Let's start with a quiver which is an orientation of a simply-laced Dynkin diagram. That is to say, an orientation of one of the following graphs:

A_n:

D_n:

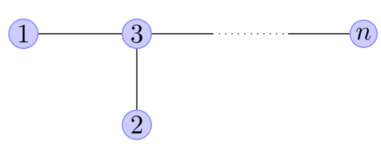

E_8:

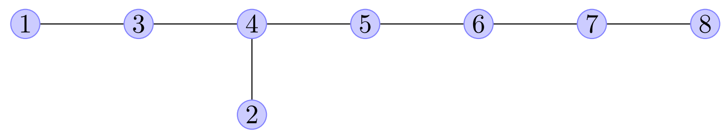

Or E_6, E_7 obtained by removing vertices 7 and 8, or just vertex 8.

A Dynkin diagram with n vertices determines a *root system* Φ, a collection of vectors ("roots") in $\mathbb{R}^n$ with remarkable properties. For our purposes, it is convenient to fix the simple root associated to vertex i of the diagram to be e_i. The root system then divides into the *positive roots*, a collection of non-negative linear combinations of the simple roots, and *negative roots*, their negatives. If you are not comfortable with root systems, that won't be a problem. It will be useful to know that in type A_n, the positive roots are of the form $e_i + e_{i+1} + \cdots + e_j$ for $i \leq j$.

Something very special about the quivers we get by orienting a Dynkin diagram is the following:

Theorem 1.1 (Gabriel's theorem)**.** *The simply-laced Dynkin quivers are exactly the quivers having only a finite number of indecomposable representations. There is one for each positive root of the corresponding root system. For γ a positive root, the generic representation of dimension γ is the indecomposable corresponding to γ, which we denote M_γ.*

We can see that our analysis A_2 was consistent with Gabriel's theorem, in that we found three different indecomposable representations. Their dimensions were (1,0), (0,1), and (1,1), which are the expansions of the three positive roots in the basis of simple roots, and they were the generic representations of their dimension vectors.

Consider $\mathrm{Rep}_Q(\mathbf{d})$, the space of representations of Q of dimension vector $\mathbf{d}$. (Recall that this is the affine space of entries of the matrices that define the representation.) Within this space, points correspond to representations, which can be grouped into different isomorphism classes. In the example with Q of type A_2, $\mathrm{Rep}_Q((1,1))$ is just $\mathbb{C}$. Now $\mathrm{Rep}_Q((1,1))$ consists of two parts, the part corresponding to $S_1 \oplus S_2$ (the origin) and the part corresponding to $M_{(1,1)}$ (everything else). If $\mathbf{d}$ is bigger, the decomposition of $\mathrm{Rep}_Q(\mathbf{d})$ will be more complicated. For example, in $\mathrm{Rep}_Q((2,2))$, we have the generic region corresponding to $M_{1,1} \oplus M_{1,1}$, we have the region corresponding to $S_1^2 \oplus S_2^2$ (the origin), but we also have a region corresponding to $M_{1,1} \oplus S_1 \oplus S_2$.

Because Dynkin quivers have only a finite number of indecomposable representations by Gabriel's theorem, there are only a finite number of isomorphism classes of representations in any particular dimension vector $\mathbf{d}$ (since the isomorphism class can be described by giving the indecomposable summands and multiplicities). This is what guarantees that there will be a "typical" behaviour: in order to account for the entire space of the representations of dimension $\mathbf{d}$, there must be at least one of the isomorphism classes which is full-dimensional, and the fact that each isomorphism class is a reasonable kind of set (technically, it is "constructible," the intersection of a Zariski open set with a Zariski closed set) there can't be more than one that is top-dimensional, and, also, each has a well-defined codimension.

It turns out there is a very algebraic way to understand the codimension of an isomorphism class in the space of representations of dimension $\mathbf{d}$.

Theorem 1.2. *Let X be a representation of dimension vector $\mathbf{d}$ of a quiver Q (not necessarily Dynkin). The codimension of the isomorphism class of X in $\mathrm{Rep}_Q(\mathbf{d})$ is $\dim \mathrm{Ext}^1(X, X)$.*

I am going to treat Gabriel's theorem and the previous theorem as black boxes; we're going to see that we can understand a lot about the geometry of the generic representation fan for Dynkin quivers using just these two facts.

Corollary 1.3. *For Q a Dynkin quiver, and any indecomposable representation M_α, we have $\mathrm{Ext}^1(M_\alpha, M_\alpha) = 0$.*

Proof. This follows from the fact that in the space of representations of dimension vector α, M_α is the generic representation, which was part of Gabriel's theorem. □

Corollary 1.4. *For any dimension vector $\mathbf{d}$ for a Dynkin quiver Q, there is a unique way to write it as a sum of positive roots $\mathbf{d} = \beta_1 + \cdots + \beta_r$ such that $\mathrm{Ext}^1(M_{\beta_i}, M_{\beta_j}) = 0$ for $1 \leq i, j \leq r$.*

Proof. This way of writing $\mathbf{d}$ must describe the generic representation of dimension vector d, which is unique. □

To talk about the generic decomposition fan for an arbitrary Dynkin quiver, I'd better introduce a little terminology about cones and fans. A *cone* in a real vector space V is (for me) always the set of points of the form $c_1v_1 + \cdots + c_rv_r$, where $v_1, \ldots, v_r \in V$, and $c_1, \ldots, c_r \in \mathbb{R}_{\geq 0}$. We say that a point is in the *relative interior* of a cone if it is in the cone and doesn't lie on any of its proper faces. A cone is called "simplicial" if it has as many generators as its dimension. (Every cone has at least that many generators, but some need more.)

A *fan* is a collection of cones such that any face of a cone in the fan is also in the fan, and the intersection of two cones is a face of each. (Don't worry too much about the technical definition of fan; for our purposes, it's really good enough to think that it's a collection of cones whose relative interiors partition the region covered by the fan.)

Consider the fan Σ_Q defined as follows:

- The rays are generated by the positive roots for Q.
- The cones are generated by subsets of the positive roots for Q corresponding to representations which have no extensions.

Corollary 1.5. *The cones of Σ_Q are all simplicial.*

Proof. If $\sigma \in \Sigma_Q$ were non-simplicial, there would be a linear dependence among its generators. This can be massaged into the statement that

$$a_1\beta_1 + \cdots + a_r\beta_r = b_1\gamma_1 + \cdots + b_s\gamma_s$$

for some $a_i, b_i \in \mathbb{Z}_{>0}$, and $\beta_1, \ldots, \beta_r, \gamma_1, \ldots, \gamma_s$ are all distinct positive roots corresponding to generators of σ. But this would show that $a_1\beta_1 + \cdots + a_r\beta_r$ admits two different generic representations, which is impossible. □

Corollary 1.6. *Σ_Q is the generic decomposition fan for Q. That is to say, for any cone σ, all the lattice points $\mathbf{d}$ in the relative interior of σ have generic decompositions which use indecomposable representations corresponding to all the rays of σ, and any dimension vector $\mathbf{d}$ lies in the relative interior of exactly one cone of Σ_Q.*

Proof. By what we have already seen, for any $\mathbf{d}$, there is a way to write it as a positive $\mathbb{Z}$-linear combination of the roots generating some cone of Σ_Q. And if $\mathbf{d}$ lies in some cone σ of Σ_Q, its generic decomposition must be given by a sum of the roots corresponding to the rays of σ. Otherwise, we would be able to find two different generic decompositions for $\mathbf{d}$ (or some multiple of $\mathbf{d}$). This also shows that the relative interiors of the cones don't overlap, so the cones overlap only on boundaries. The intersection of two cones will necessarily be the cone generated by the intersection of their two sets of generating roots. $\square$

1.3. **Beyond hereditary type.** While the theory of Kac [**?**] extends the generic decomposition fan to quivers beyond Dynkin type, in the interests of time, I am going to skip that, and ask what we can do when the quiver is not hereditary. Consider the following quiver with relations:

$$\begin{tikzcd} 1 \arrow[r, bend left, "\alpha"] & 2 \arrow[l, bend left, "\beta"] \end{tikzcd}$$

Here the dotted lines mean that I insist that the corresponding compositions of arrows be zero. Consider the representations X of dimension vector $(1,1)$ of this quiver with relations:

$$\begin{tikzcd} \mathbb{C} \arrow[r, bend left, "{[c_\alpha]}"] & \mathbb{C} \arrow[l, bend left, "{[c_\beta]}"] \end{tikzcd}$$

Now, c_α and c_β are just scalars, and the relations mean that $c_\alpha c_\beta = 0$. The space of all representations of dimension vector (1,1) for this quiver is no longer an affine space: it is cut out by equations coming from the relations. (This would still be true if the relations were more complicated.) The space of representations forms the two coordinate axes (only).

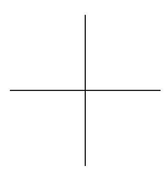

At the origin, we have $S_1 \oplus S_2$, but on the two axes, we have two different isomorphism classes of representations. In particular, there is no generic choice of representation!

There is a beautiful way to solve this problem, due to Derksen and Fei [**?**], which was one of the forerunners of τ-tilting theory as developed by Adachi, Iyama, and Reiten [**?**]. Rather than pick a dimension vector, and look for a generic behaviour there, we will pick a generic projective presentation.

Let $P(1), \ldots, P(n)$ be the indecomposable projective representations of (Q, I). We pick an n-tuple $(g_1, \ldots, g_n) \in \mathbb{Z}$, which we will call a g-vector. Define

$$P^0 = \bigoplus_{i:g_i>0} P(i)^{g_i} \qquad P^{-1} = \bigoplus_{i:g_i<0} P(i)^{-g_i}$$

Now observe that the space of morphisms from P^{-1} to P^0 is just an affine space (since $\mathrm{Hom}(X, Y)$ is always a vector space), and we can ask about the behaviour of a generic presentation. (This requires that we introduce a suitable category of these presentations. This is the "homotopy category of two-term complexes over proj Q". More details will follow!)

I am running out of time, so I want to explain only how this recovers the generic decomposition fan (plus a bit more) in the Dynkin case.

The dimension vectors of the indecomposable projective representations form a basis for $\mathbb{Z}^n$. If we have any dimension vector d, we can write it in the basis of projectives as $g \in \mathbb{Z}^n$. The generic presentation at such a g-vector is a projective resolution of the generic module with dimension vector d.

The picture of the g-fan for $1 \leftarrow 2$ is as follows:

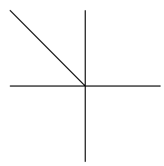

Exercises.

(1) Find the generic decomposition fan for the quivers

$$1 \xrightarrow{\alpha} 2 \xrightarrow{\beta} 3 \quad \text{and} \quad 1 \xrightarrow{\alpha} 2 \xleftarrow{\beta} 3$$

If this is too straightforward, find the generic decomposition fan for $1 \leftarrow 2 \cdots \leftarrow n$.

(2) Locate the generic decomposition fan for $1 \leftarrow 2$ in the g-fan for $1 \leftarrow 2$. Hint: it's not in the most obvious place.

(3) Find equations which cut out the complement of the region where the representation of dimension (d_1, d_2) with $d_2 < d_1$ is injective. Describe what happens generically when $d_2 > d_1$.

(4) Show that there is no generic representation with dimesion (1,1) of the Kronecker quiver:

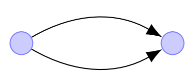

Why does the analysis we gave above of Dynkin quivers not apply? (Beyond the obvious fact that this is not a Dynkin quiver.)

2. The u-equations

2.1. **The g-fan for linearly oriented A_3.** As we discussed last time, for quivers, the rays of the generic decomposition fan turn into rays for the g-fan, once we express the dimension vectors in terms of the basis of projectives. Here is the picture for $1 \leftarrow 2 \leftarrow 3$, where we represent the 3d picture by taking a 2d slice of positive orthant and projecting into the slice by lines through the origin. The lefthand side shows the dimension vectors; the righthand side shows the g-vectors, where we used that the dimensions are $\dim(P(1)) = (1, 0, 0), \dim(P(2)) = (1, 1, 0), \dim(P(3)) = (1, 1, 1)$.

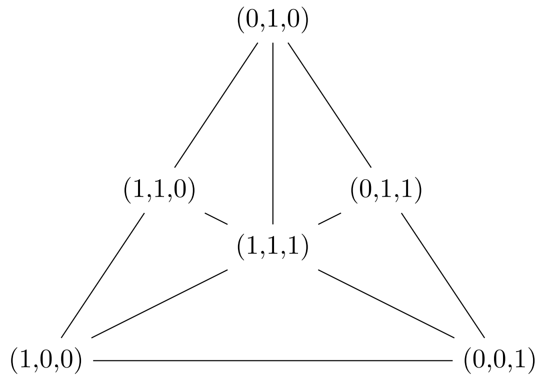

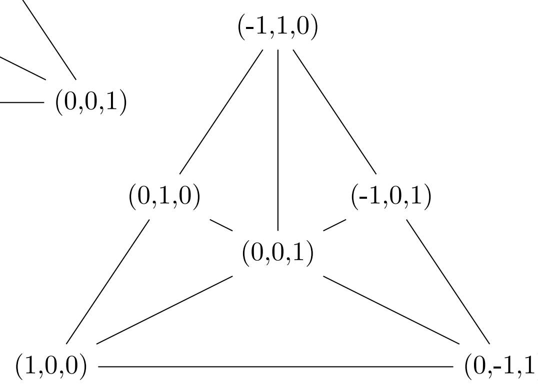

However, that is not the whole of the g-fan: we also add in an additional n rays (here $n = 3$) for the presentations $P(i) \to 0$. The ray $P(i) \to 0$ is compatible exactly with the presentations which don't have any support over vertex i. The corresponding g-vector is $-e_i$. The full g-fan for Dynkin quivers (or for any τ-finite fan) covers the entire space, so I'm just going to represent it combinatorially.

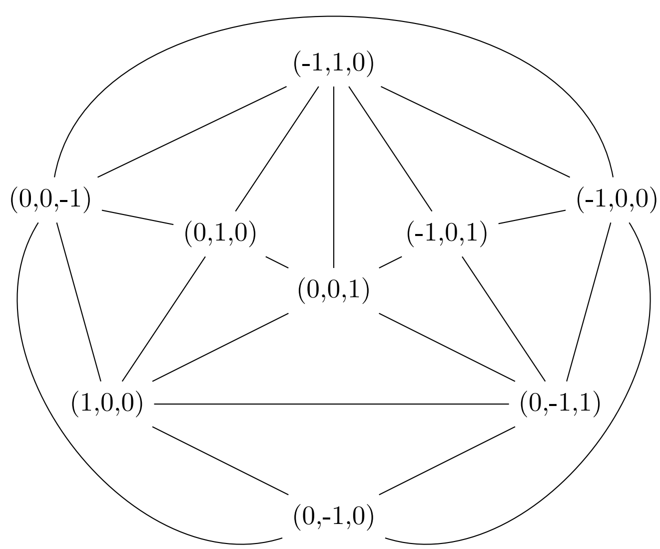

2.2. **Combinatorics of the A_n g-vector fan.** In Daniel's talk today, he explained that if we start with an $(n+3)$-gon, and we fix a triangulation T for it, there is an associated algebra Λ. If we choose T so that each triangle has an edge on the boundary, then Λ is the path algebra of an A_n quiver; otherwise, Λ will be a gentle algebra of finite representation type.

Then, the rays of the g-fan correspond to diagonals of the polygon, and pairs of rays lie in the same cone if the corresponding diagonals don't cross. I'm going to need to make that statement more precise in a moment.

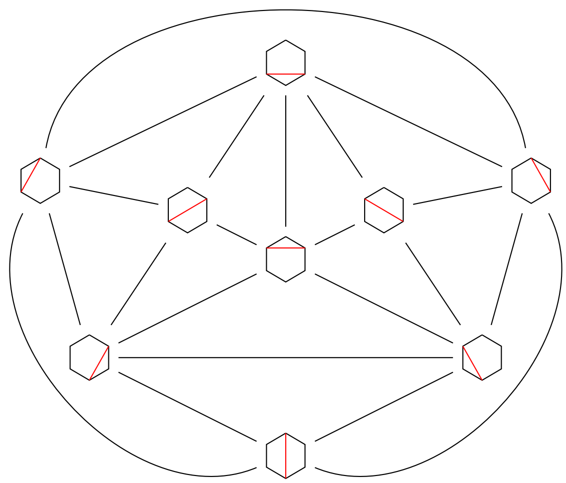

2.3. **u-equations of Koba and Nielsen.** We introduce variables u_{ij} indexed by $ij \in D$. We also introduce an equation (which we will call the ij equation) for each diagonal $ij \in D$:

$$u_{ij} + \prod_{kl \in D,\, kl \text{ crosses } ij} u_{kl} = 1.$$

(Two diagonals are considered to cross only if they intersect in their interiors; having a common endpoint does not count.)

These equations go back to work of physicists Koba and Nielsen in 1969 [KN1, KN2, KN3]. Koba and Nielsen were particularly interested in the real, non-negative solutions of this system of equations. Let's call all the complex solutions $\mathcal{U}$, and the non-negative real solutions $\mathcal{U}_{\geq 0}$ and see what we can say about them.

First of all, if all the u_{ij} are non-negative, they must also all be at most 1: for u_{ij} this follows from the ij equation.

Define a "face" of $\mathcal{U}_{\geq 0}$ to be the region defined by some subset of the $u_{ij} = 0$. ($\mathcal{U}_{\geq 0}$ is definitely not a polytope, but it is supposed to "feel" like a polytope.) Order the faces of $\mathcal{U}_{\geq 0}$ by inclusion.

Theorem 2.1. *The poset of faces of $\mathcal{U}_{\geq 0}$ is anti-isomorphic to the poset of cones of Σ_Q.*

Proof. Which collections of diagonals can be zero at once? Well, if u_{ij} is zero, then any diagonal crossing ij must be equal to 1 by the ij equation and the fact, already shown, that all the $u_{kl} \leq 1$. If we take a maximal collection of non-crossing diagonals (i.e., a triangulation), then we can set all the u_{ij} in the triangulation to zero, and all the others to 1. One easily sees that this satisfies all the u-equations. (Exercise: check this.) So, the faces necessarily correspond to partial triangulations, i.e., for every cone of the g-fan, there is a corresponding face, and these are all of them. The faces corresponding to different cones of the g-fan are different, because the face corresponding to σ contains the zero-dimensional faces corresponding to maximal cones containing σ (and no others). So we have a bijection. □

How reasonable is it to call these things "faces" of $\mathcal{U}_{\geq 0}$? It turns out it's pretty reasonable.

There is a key fact about solving the u-equations which, for the moment, I am going to assume: $\mathcal{U}$ defines an n-dimensional family of solutions, and $\mathcal{U}_{\geq 0}$ is real n-dimensional. Under this assumption we can establish:

Theorem 2.2. *The face of $\mathcal{U}_{\geq 0}$ corresponding to a d-dimensional cone σ of Σ_Q is of real dimension $n - d$.*

Proof. (This proof depends on the assumption that $\mathcal{U}_{\geq 0}$ has has the correct dimension, which we will justify later in today's talk.) By induction, it's enough to check that $\mathcal{U}_{\geq 0} \cap \{u_{ij} = 0\}$ is $(n-1)$-dimensional. As already observed, the ij equation implies that all the diagonals kl that cross ij must have $u_{kl} = 1$. The remaining equations are simply the u-equations for the polygons on the two sides of diagonal ij. These are a $(k+1)$-gon and an $(n+3-k+1)$-gon. (We split up the edges of the $n+3$-gon between the two polygons, but each also gets one new edge from the dividing diagonal.) Thus, the corresponding $\mathcal{U}$-spaces are of dimension $k-2$ and $n-k+1$, whose sum is $n-1$. Therefore (under our assumption that u-equations have the predicted dimension of solutions) we see that the region where u_{ij} is zero is $(n-1)$-dimensional, as desired. □

The purpose of Koba and Nielsen was to define something called the *tree string integral*; their idea for how to do it was to write it as an integral over $\mathcal{U}_{\geq 0}$. This integral is a function of "Mandelstam variables" which correspond to u-variables. If a Mandelstam variable s_{ij} goes to zero, then the main contribution to the integral comes from wery near the locus where $u_{ij} = 0$. The desired feature is that this should exhibit "factorization," splitting apart into two smaller instances of the same problem. So the fact that when $\mathcal{U}_{\geq 0} \cap \{u_{ij} = 0\}$ splits as a product of two smaller instances of totally non-negative $\mathcal{U}$ spaces(corresponding to the polygons on either side of ij, is not just a curiosity, but is in fact the whole point.

I should clarify that Koba and Nielsen also have a way to solve their u-equations, which involves interpreting u_{ij} as a cross-ratio of points on the projective line.

2.4. F-polynomials.

F-polynomials were first introduced by Derksen–Weyman–Zelevinsky [DWZ]. The F-polynomial of a representation is a generating function that "counts" its submodules. Since we like to work over infinite fields (and in fact, we will prefer to work over $\mathbb{C}$), we cannot literally count them; counting must be replaced by taking a suitable Euler characteristic.

Formally, for Q a quiver with n vertices numbered 1 to n, we introduce formal variables $y_1, \ldots, y_n$. For d a dimension vector, we write y^d for $y_1^{d_1} y_2^{d_2} \ldots y_n^{d_n}$.

For M a representation of Q over $\mathbb{C}$, we define

$$F_M = \sum_{d \in \mathbb{Z}_{\geq 0}} \chi \operatorname{Gr}(M,d) y^d$$

The quiver Grassmannian $\operatorname{Gr}(M,d)$ is a projective subvariety of usual Grassmannians $\prod_{i=1}^n \operatorname{Gr}(M_i, d_i)$, consisting of those choices of subspaces which actually define a subrepresentation, i.e., such that the maps of M restrict to maps between the chosen subspaces, and χ is the Euler characteristic.

Note that the sum is actually finite, because $\operatorname{Gr}(M,d)$ is only non-empty if d is coordinate-wise less than or equal to $\dim M$.

While the Euler characteristics may seem off-putting, in many cases the quiver Grassmannians will just be points, with Euler characteristic 1. (It's also useful to know that the Euler characteristic of $\mathbb{P}^n$ is $n+1$.)

For the quiver $1 \leftarrow 2 \leftarrow \cdots \leftarrow n$, it is easy to determine the F-polynomials. The F-polynomial of the indecomposable representation supported on vertices $i, i+1, \dots, j$ is $1 + y_i + y_i y_{i+1} + \cdots + y_i y_{i+1} \dots y_j$.

2.5. **The type A_n g-fan in more detail.** Fix n a positive integer, at least 1, and consider a convex $(n+3)$-gon R. Write D for the diagonals of R. Edges do not count as diagonals. The number of diagonals $|D| = n(n+3)/2$.

Pick a triangulation T of R. Our triangulation might look like this:

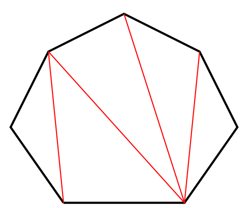

Put a vertex on each of the edges of the triangulation, and connect two vertices by an arrow if they lie in the same triangle. Orient the arrow so it is moving clockwise around its own triangle. (This is the construction Daniel sketched in the problem session on Day 1.)

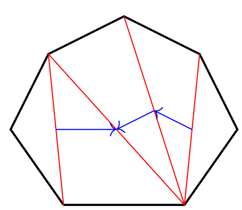

If we choose the triangulation so that triangle has an edge on the boundary, the resulting quiver will be of type A_n, and we let $\Lambda = \mathbb{C}Q$. If not, the quiver will have some 3-cycles. In this case, let Λ be the path algebra $\mathbb{C}Q$ modulo the quadratic relations consisting of any two consecutive arrows in a 3-cycle. (For my purposes, it would be good enough to stick to A_n quivers, or even just a single A_n quiver, so feel free to ignore this if it's confusing.)

We will refer to as a *laminate* a curve C which starts and ends on the boundary of R, but not at one of the vertices. We also exclude curves which enter and leave by the same edge or by adjacent edges. A laminate is really an equivalence class of curves: two laminates with endpoints on the same pair of segments are equivalent.

(If we were dealing with a surface with non-trivial topology, we would also need to require that they be homotopic.)

Our goal in this subsection is to make explicit bijections among:

- rays of the g-fan for Λ
- g-vectors for Λ
- indecomposable representations of Λ plus $\{P(i) \to 0 \mid 1 \leq i \leq n\}$
- diagonals of R
- laminates of R

The first three things are equivalent by the theory of g-fans, as discussed by Osamu on Days 1 and 2. The latter two also have a simple bijection: rotate the endpoints of the diagonal slightly counter-clockwise to obtain the corresponding laminate.

So the real goal of what I am presenting in this subsection is to make the connection between the combinatorics of diagonals/laminates and g-vectors/representations.

A diagonal which is not in T corresponds to the indecomposable representation supported at the vertices which it crosses; this uniquely defines it; the diagonals in T correspond to the presentations $P(i) \to 0$.

To see the g-vector combinatorially, it is better to work with laminates. The *shear coordinates* of a laminate C with respect to the triangulation T are defined as follows. Number the edges of T arbitrarily $T_1, \ldots, T_n$. Then $\mathrm{sh}_T(C)_i$ is determined based on how C passes through the quadrilateral whose diagonal is T_i.

If it doesn't cross T_i at all, then $\mathrm{sh}_T(C)_i = 0$. If the three arcs it passes through form an S shape, the $\mathrm{sh}_T(C)_i = 1$. If they form a Z shape, then $\mathrm{sh}_T(C)_i = -1$. If they form an arrow, the shear coordinate is zero. The lefthand image is an example of a -1 shear coordinate; the second image is an example of a $+1$ shear co-ordinate, and the last image is an example of a zero shear coordinate.

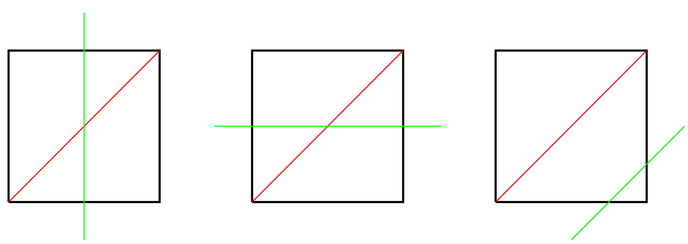

The g-vector corresponding to a laminate is given by its shear coordinates. (This can be deduced either via a detour through categorification of cluster algebras [DK, FT], or combinatorially.)

2.6. **Extension to arbitrary surfaces without punctures.** As Daniel explained, to a triangulated surface without punctures, there is a corresponding gentle algebra [**?**].

The story above applies! Indecomposable representations (plus $P(i) \to 0$) correspond to laminates ("string modules") plus non-contractible closed curves ("band modules"). In both cases, the g-vector is calculated as above. The dimension is determined by crossings of the triangulation by the corresponding "diagonal": this is a curve between vertices obtained by moving the endpoints clockwise to vertices.

Strings with self-crossings and bands are not τ-rigid, so they don't contribute to the g-vector fan.

2.7. **Solving the u-equations of Koba and Nielsen.** To solve the u-equations using representation theory, I am going to pick a particular triangulation T, namely, the one which connects one vertex to all the other vertices. I'm going to number the edges of T in clockwise order. The quiver corresponding to this triangulation is $1 \leftarrow 2 \leftarrow \cdots \leftarrow n$.

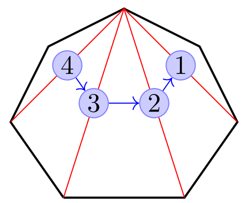

The g-vectors for this triangulation are $e_1, \ldots, e_n; -e_1, \ldots, -e_n; e_j - e_i$ for $j > i$. Below, we see examples for e_2, $-e_3$, and $e_4 - e_2$.

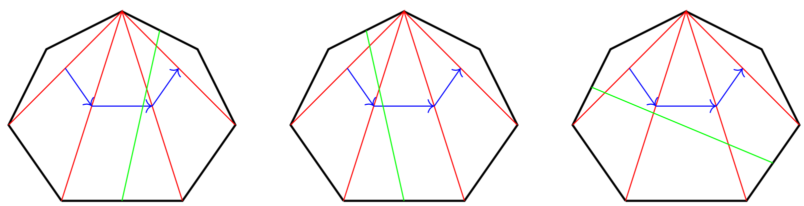

I'm going to introduce a certain way of organizing the g-vectors into an array. Those who are familiar with AR theory will recognize it as the AR quiver; if that's not familiar, don't worry about it. In the next lecture, I will be presenting a generalization of this solution which doesn't require the use of the AR quiver.

$$\begin{array}{ccccccccc}
 & & & e_4 & & -e_1 & & & \\
 & & e_3 & & e_4 - e_1 & & -e_2 & & \\
 & e_2 & & e_3 - e_1 & & e_4 - e_2 & & -e_3 & \\
e_1 & & e_2 - e_1 & & e_3 - e_2 & & e_4 - e_3 & & -e_4
\end{array}$$

The reason to draw this picture is that, for a g-vector $e_j - e_i$, the other term in the ij equation is the product of the u's corresponding to two rectangles.

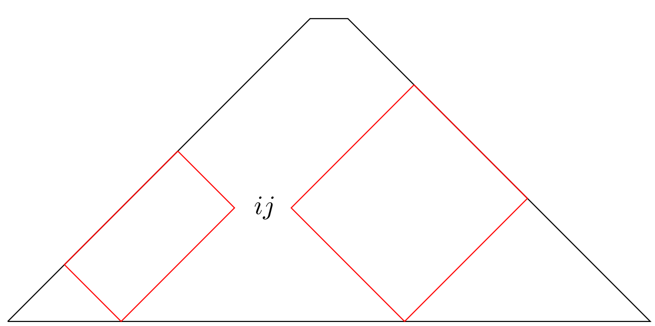

For convenience, let us declare $e_0 = e_{-1} = 0$. In place of writing F_M, I'm going to write F_g where g is the g-vector of M. Also, it's convenient to set $F_0 = 1$. Define

$$u_{e_j - e_i} = \frac{F_{e_{j-1}-e_i} F_{e_j - e_{i-1}}}{F_{e_{j-1}-e_{i-1}} F_{e_j - e_i}} \text{ for } 1 \leq j \leq n, \text{ and } 0 \leq i < j$$
$$u_{-e_i} = \frac{y_i F_{e_n - e_i}}{F_{e_n - e_{i-1}}}$$

Part of the asymmetry in the definition could be avoided if we had just declared that $F_{-e_i} = 1$, which is not unreasonable, since that g-vector doesn't correspond to an actual module.

Now, when we take the products of the u's in the two rectangles drawn above for the ij equation, they will telescope.

Lemma 2.3. *The product of the u's in the lefthand rectangle is $\frac{1}{F_{e_{j-1}-e_{i-1}}}$. The product of the u's in the righthand rectangle is $\frac{y_{i+1} \cdots y_j}{F_{e_j - e_i}}$.*

Then it is just an easy calculation, using the F's that we calculated before, to conclude:

Theorem 2.4. *The formulas given above provide a rational parametrization of an n-dimensional family of solutions to the u-equations.*

In fact, more is true: the region $\mathcal{U}_{>0}$ is parametrized by the $y_i > 0$. (It is clear that if the y_i are positive, then the u_{ij} are positive; it turns out that the converse is also true.) The parametrization by y_i's is also useful because Koba and Nielsen need a measure on $\mathcal{U}_{\geq 0}$; this turns out to be very easily described in terms of the y_i as $\bigwedge_i dy_i/y_i$.

(For those keeping track, this also justfies the assumption we made in our discussion at the beginning, that $\mathcal{U}$ was n complex dimensional and that $\mathcal{U}_{\geq 0}$ was n real dimensional.)

Exercises.

(1) Check the assertion made in the text that the ij equation is
$$u_{ij} + \prod_{kl \text{ in one of the two rectangles}} u_{kl} = 1$$
(2) Check the F-polynomial calculations for the quiver $1 \leftarrow 2 \leftarrow \cdots \leftarrow n$.
(3) Prove the assertion made in the text that the quiver resulting from our choice of triangulation will necessarily be of type A_n. Characterize the quivers that can result for other choices of triangulation.
(4) Find a representation of a quiver for which a coefficient other than 1 appears in the F-polynomial.
(5) Show that the rays in the g-fan corresponding to $P(i) \to 0$ and a minimal projective presentation of a module with support at vertex i do not span a cone in the g-fan. (Hint: think about what the generic presentation with a g-vector that is a positive linear combination of these two g-vectors would look like, and see that it is not a direct sum of these two.)

(6) (For those familiar with AR theory) Consider a different orientation of the A_n quiver, and see how the same approach still works to solve the u-equations. (A typical u_{ij} will be the ratio of four F's from a mesh of the AR quiver.)

3. u-EQUATIONS FOR FINITE REPRESENTATION TYPE ALGEBRAS

3.1. **Defining u-equations.** Let Λ be a finite-dimensional algebra of finite representation type.

Let $\mathcal{I}$ be the set $\operatorname{ind}\Lambda \cup \{P(i)[1] \mid 1 \leq i \leq n\}$. (It is also possible to view $\mathcal{I}$ as the set of indecomposable objects in $K^2(\operatorname{proj}\Lambda)$. In some respects, it's preferable to do so. But in the interests of keeping things elementary, I'm going to avoid doing so for today.)

Introduce a variable u_M for each $M \in \mathcal{I}$.

Define the compatibility degree c of $M, N \in \mathcal{I}$ by

$$c(M,N) = \dim \operatorname{Hom}(M, \tau N) + \dim \operatorname{Hom}(N, \tau N).$$

Here, we interpret $\tau P(i)[1] = I(i)$, and $\operatorname{Hom}(P(i)[1], M) = 0$ for $M \in \operatorname{mod}\Lambda$.

The equation associated to M is:

$$u_M + \prod_{N \in \mathcal{I}} u_N^{c(M,N)} = 1.$$

Again, we consider $\mathcal{U}$, the complex solutions to this system of equations, and $\mathcal{U}_{\geq 0}$, the non-negative real solutions.

Lemma 3.1. *If Λ is an algebra associated to a disk, then this recovers the Koba–Nielsen u-equations from Lecture 2.*

Proof. This is easy to see if the algebra associated to the disk is a Dynkin quiver, where $c(M,N) = \dim \operatorname{Ext}^1(M,N) + \dim \operatorname{Ext}^1(N,M)$; for the gentle case, it's possible to use the combinatorial description of $\operatorname{Hom}(M, \tau N)$ in [BD+]. $\square$

What I want to do next is very analogous to what we did in Section 2.3.

Again, we want to define "faces" of $\mathcal{U}_{\geq 0}$ to be $\mathcal{U}_{\geq 0} \cap \{u_{M_1} = \cdots = u_{M_r} = 0\}$.

Lemma 3.2. *The poset of faces of $\mathcal{U}_{\geq 0}$ is anti-isomorphic to the poset of cones in the g-fan of Λ.*

Proof. If $M \in \mathcal{I}$ is not τ-rigid, then by definition $\operatorname{Hom}(M, \tau M) \neq 0$, and $c(M,M) > 0$. This means that the M equation has two terms on the lefthand side, both having a factor of u_M. It follows that there are no solutions to $u_M = 0$ in $\mathcal{U}_{\geq 0}$, so M doesn't define a boundary. (This is not something we had to consider yesterday, because for the algebras associated to a disk, every indecomposable is τ-rigid.)

On the other hand, if $M \in \mathcal{I}$ *is* τ-rigid, then the M equation implies that $u_N = 1$ for all N which are incompatible with M, and the argument goes through the same way as for A_n. $\square$

Again, we will assume that $\mathcal{U}_{\geq 0}$ is of real dimension n (which, again, we will justify later today).

What remains of the u-equations when we set $u_M = 0$ for one M? The equations for N incompatible with M are trivially satisfied, so the other remaining equations are for $N \in \mathcal{I}$ with $c(M,N) = 0$:

$$u_N + \prod_{K\in\mathcal{I},c(M,K)=0} u_K^{c(N,K)} = 1$$

The bad news is that the collection of objects from $\mathcal{I}$ compatible with M, other than M itself, does not form a particularly nice collection of objects. However, Jasso [J] has shown that there is a finite-dimensional algebra of rank $n-1$, which I'll denote Λ_M, with the property that $\mathcal{I}_M$ (i.e., the set $\mathcal{I}$ for Λ_M) is in natural bijection with the objects of $\mathcal{I}$ which are compatible with M, other than M. Further, writing N_M for the object of $\operatorname{mod}\Lambda_M$ corresponding to N, we have that $c(N,K) = c_M(N_M, K_M)$. This shows that boundaries corresponding to a singlue $u_M = 0$ are of dimension $(n-1)$, and the rest follows by induction.

3.2. **Quick recall of AR sequences.** A reference for this is [ASS, Chapter 4]. A *split* short exact sequence is one which is isomorphic to the obvious short exact sequence $0 \to X \to X \oplus Y \to Y \to 0$, with the maps being inclusion and projection. In a split sequence, any morphism from Z to Y, factors through the middle term of the sequence. And in fact, this is equivalent to the sequence being split. Indeed, if the identity map from Y to itself factors through Z, then the sequence is split.

Auslander and Reiten defined what it means for a short exact sequence $0 \to X \to E \to Y \to 0$ to be almost split. Since we want almost split sequences not to be actually split, the identity map from Y to itself mustn't factor through E. And, further, if $f : Z \to Y$ is a *retraction*, i.e., if there exists a function $g : Y \to Z$ with $f \circ g = 1_Y$, then it must not be the case that f factors through E, since this would also give us a way to split the sequence.

A sequence as above is called *almost split* if any map from Z to Y which isn't a retraction, necessarily factors through E. It turns out that the dual property (regarding maps from X to Z factoring through E) also holds for almost split sequences.

It turns out that for X any non-projective indecomposable module, there is a unique almost split short exact sequence ending with X; it begins with τX. (One can also start from Y any non-injective indecomposable module; then there is a unique almost split short exact sequence beginning with Y.)

3.3. **Solving the u-equations.** The solution is as follows. For X a non-projective module, let

$$0 \to \tau X \to E \to X \to 0$$

be the Auslander–Reiten sequence ending at X. For X projective, for convenience, we define E to be the radical of X, and τX to be 0.

For X not shifted projective, define

$$u_X = \frac{F_E}{F_{\tau X} F_X}.$$

For $P(i)[1]$, define

$$u_{P(i)[1]} = \frac{y_i F_{I_i/S_i}}{F_{I_i}}$$

For $Q = 1 \leftarrow 2$, these solutions (which agree with what we saw in Lecture 2) are:

$$u_{P(2)} = \frac{1+y_1}{1+y_1+y_1y_2} \qquad u_{P(1)[1]} = \frac{y_1(1+y_2)}{1+y_1+y_1y_2}$$

$$u_{S(1)} = \frac{1}{y_1+1} \qquad u_{S(2)} = \frac{1+y_1+y_1y_2}{(1+y_1)(1+y_2)} \qquad u_{P(2)[1]} = \frac{y_2}{1+y_2}$$

In the type A case, we said that the M equation was of the form u_M+ products over a pair of rectangles in the AR quiver. These "two rectangles" are given by $\mathrm{Hom}(M, \tau N) \neq 0$ and $\mathrm{Hom}(N, \tau M) \neq 0$. For those who know the term, the rectangles are really "hammocks."

To analyze them we have the following lemmas:

Lemma 3.3.

$$\prod_N u_N^{\dim \mathrm{Hom}(N,\tau M)} = \frac{1}{F_{\tau M}}$$

Proof. The proof is by a telescoping argument. The idea is that if we have some u_N that might appear, we take the AR sequence beginning in N:

$$0 \to N \to F \to \tau^{-1}N \to 0$$

The point is that, in the product we are interested in, F_N will show up as a result of terms for u_N, $u_{\tau^{-1}N}$ and u_{F_i}, where F_i are the indecomposable summands of F. The reason that the total number of occurrences of F_N adds up to zero is that

$$0 \to \mathrm{Hom}(\tau^{-1}N, \tau M) \to \mathrm{Hom}(F, \tau M) \to \mathrm{Hom}(N, \tau M)$$

is exact on the left, and by almost splitness, it is exact on the right except if $N = \tau M$, which is what accounts for the $F_{\tau M}$ that appears in the answer. □

Lemma 3.4.

$$\prod_N u_N^{\dim \mathrm{Hom}(M,\tau N)} = \frac{y^{\dim M}}{F_M}.$$

Proof. The monomial in y appears from the $u_{P(i)[1]}$ terms. The telescoping argument works the same way and accounts for the $1/F_M$. □

Theorem 3.5. *The u_M defined above do solve the u-equations.*

Proof. Using the above two lemmas, we need to show:

$$u_M + \frac{y^{\dim M}}{F_{\tau M} F_M} = 1,$$

which is to say

$$F_M F_{\tau M} = F_E + y^{\dim M}$$

Let's notice that this is plausible. E is the almost split extension of M by τM. If it were the split extension, it is well-known that $F_{M \oplus \tau M} = F_M F_{\tau M}$. So we are saying, because the extension is almost split, there is one term from $F_M F_{\tau M}$ that is missing. Happily, this formula was proved by Domínguez and Geiss [DG]! □

Note that we have showed that the above formulas give solutions to the u-equations, but we haven't showed that they are the only ones. To be more precise, the above formulas give us a rational parametrization of one component of $\mathcal{U}$. However, we know it is the component that contains $\mathbb{U}_{\geq 0}$, which means it's the component we want. Further, we could conjecture:

Conjecture 3.6. *$\mathcal{U}$ has only one component.*

I am not very committed to this conjecture. It is known in the Koba–Nielsen case discussed last time, and somewhat more generally, by [AHL]. I am committed to the idea that we only care about this component, but in fact there is a replacement for the set of u-equations which cuts out only the component we like.

3.4. **$\widehat{F}$-equations.** For M in $\operatorname{ind} \operatorname{mod} \Lambda$, define

$$\widehat{F}_M = \sum_{\mathbf{d}} \chi \operatorname{Gr}(M, \mathbf{d}) \prod_{N \in \mathcal{I}} u_N^{\dim \operatorname{Hom}_\Lambda(N,M) - \langle g_N, \mathbf{d} \rangle}$$

This is the result of carrying out a monomial change of variables in the F-polynomial of M, followed by multiplication by an overall monomial. (This is reminiscent of what happens when passing from principal coefficients to arbitrary coefficients in a cluster algebra. In that case, the overall scaling factor would be the tropical evaluation of F, i.e., for each u_M, we multiply by the smallest power of u_M to that the result is polynomial in y_M. A result of Fei [F1, F2] shows that the two coincide with that evaluation, if all indecomposables are rigid, as in the Koba–Nielsen case.)

For $1 \leftarrow 2$, the $\widehat{F}$ equations are:

$$\widehat{F}_{S(1)} = u_{S(1)} + u_{S(2)} u_{P(1)[1]} \qquad \widehat{F}_{S_2} = u_{P(2)} u_{S(2)} + u_{P(2)[1]}$$
$$\widehat{F}_{P(2)} = u_{S(1)} u_{P(2)} + u_{P(2)} u_{S(2)} u_{P(1)[1]} + u_{P(1)[1]} u_{P(2)[1]}$$

Now, the equations which replace the u-equations are the equations that $\widehat{F}_M = 1$. (Note while there are as many u-equations as variables, there are n fewer $\widehat{F}$ equations than there are variables, which is consistent with the fact that the solutions wind up being n-dimensional.)

Lemma 3.7. *If $\{u_N\}$ satisfy the $\widehat{F} = 1$ equations, then they satisfy the u-equations.*

Proof. Suppose for convenience that X is not projective or shifted projective, and suppose the AR-sequence ending at X is $0 \to \tau X \to E \to X \to 0$. Domìnguez–Geiss [DG] tells us that $F_X F_{\tau X} = F_E + u^{\dim X}$. We can lift this to $\widehat{F}$'s by doing the monomial change of variables (always the same) and taking account of the global factors (different for each $\widehat{F}$). The Domínguez–Geiss formula lifts to:

$$\widehat{F}_X \widehat{F}_{\tau X} = u_X \widehat{F}_E + \prod_{N \in \mathcal{I}} u_N^{c(N,X)}$$

Now, if all the $\widehat{F}_X = 1$, we recover the u-equation! □

In the $1 \leftarrow 2$ example, the above equation is:

$$\widehat{F}_{S(1)} \widehat{F}_{S(2)} = u_{S(2)} \widehat{F}_{P(1)} + u_{S(1)} u_{P(2)[1]}$$

Theorem 3.8. *The formulas we gave in Section 3.3 also solve the $\widehat{F}$-equations.*

Proof. The proof again is a telescoping argument. When we carry out the substitution, each term of $\widehat{F}_M$ turns into the corresponding term of F_M, divided by an overall factor of F_M. The sum of all the terms is therefore $F_M / F_M = 1$. □

I begin by recalling the setup from last time.

Λ is a finite dimensional algebra of finite representation type. $\mathcal{I}$ is our indexing set, which we can think about either as indecomposable Λ modules plus $P(i)[1]$, or as indecomposable complexes in the homotopy category of two-term complexes. We introduced a variable u_M for each $M \in \mathcal{I}$, and a compatibility degree $c(M,N) = \dim \operatorname{Hom}(M, \tau N) + \dim \operatorname{Hom}(N, \tau M)$.

Recall from last time: we had defined two systems of equations: the u-equations:

$$u_M + \prod_{N \in \mathcal{I}} u_N^{c(M,N)} = 1$$

And the $\widehat{F} = 1$ equations, where $\widehat{F}_M$ was a lift of F_M in the variables u_M.

$$\widehat{F}_M = \sum_{\mathbf{d}} \chi \operatorname{Gr}(M, \mathbf{d}) \prod_{N \in \mathcal{I}} u_N^{\dim \operatorname{Hom}_\Lambda(N,M) - \langle g_N, \mathbf{d} \rangle}$$

(One thing one might notice, though I didn't say it, is that there are $|\mathcal{I}|$ many u-equations, but only $|\mathcal{I}| - n$ many $\widehat{F}$-equations; and the latter is the actual codimension of $\mathcal{U}$.) In the case $1 \leftarrow 2$, we had:

$$\widehat{F}_{S(1)} = u_{S(1)} + u_{S(2)} u_{P(1)[1]} \qquad \widehat{F}_{S_2} = u_{P(2)} u_{S(2)} + u_{P(2)[1]}$$
$$\widehat{F}_{P(2)} = u_{S(1)} u_{P(2)} + u_{P(2)} u_{S(2)} u_{P(1)[1]} + u_{P(1)[1]} u_{P(2)[1]}$$

The rational functions which solve the u-equations, and the $\hat{F}$-equations, are:

$$u_X = \frac{F_E}{F_{\tau X} F_X} \qquad (X \text{ not shifted projective})$$
$$u_{P(i)[1]} = \frac{y_i F_{I_i/S_i}}{F_{I_i}}$$

In the $1 \leftarrow 2$ case, we had the following rational parametrization of solutions:

$$u_{P(2)} = \frac{1 + y_1}{1 + y_1 + y_1 y_2} \qquad u_{P(1)[1]} = \frac{y_1(1+y_2)}{1 + y_1 + y_1 y_2}$$
$$u_{S(1)} = \frac{1}{y_1 + 1} \qquad u_{S(2)} = \frac{1 + y_1 + y_1 y_2}{(1+y_1)(1+y_2)} \qquad u_{P(2)[1]} = \frac{y_2}{1+y_2}$$

Also recall that we had $\mathcal{U}_{\geq 0} \subseteq \overline{\mathcal{U}_{\geq 0}} \subseteq \mathcal{F} \subseteq \mathcal{U}$. (Here I am introducing new notation: $\overline{\mathcal{U}_{\geq 0}}$ is the component of $\mathcal{U}$ containing $\mathcal{U}_{\geq 0}$, and $\mathcal{F}$ is the locus where we have $\widehat{F}_M = 1$ for all $M \in \operatorname{ind} \Lambda$.

3.5. **Torus action and irreducibility.** Consider the polynomial ring $\mathbb{C}[u_M \mid M \in \mathcal{I}]$. This ring admits a natural $\mathbb{Z}^{|\mathcal{I}|}$ grading. There is a particular grading that will be important for us. It is somewhat less refined. The grading group is $\mathbb{Z}^{|\operatorname{ind} \operatorname{mod} \Lambda|}$. The advantage of it is that the $\widehat{F}$ polynomials are homogeneous with respect to that grading. (And the grading itself can be deduced from the construction of the $\widehat{F}$ polynomials as the result of carrying out a monomial substitution starting with polynomials in only n variables.) Further, the degrees of the $\widehat{F}$ polynomials form a basis for the grading group.

The existence of a grading by $\mathbb{Z}^{|\operatorname{ind} \operatorname{mod} \Lambda|}$ on $\mathbb{C}[u_M]$ gives rise to an action of $T = (\mathbb{C}^*)^{|\operatorname{ind} \operatorname{mod} \Lambda|}$ on $\operatorname{Spec} \mathbb{C}[u_M] \simeq \mathbb{C}^{|\mathcal{I}|}$. Because the $\widehat{F}_M$ and u_M are homogeneous with respect to the grading, we can remove the locus where they are zero, and we still have a T action on $X = \mathbb{C}^{\mathcal{I}} \setminus V(\prod_{M \in \operatorname{ind} \operatorname{mod} \Lambda} u_M \prod_{M \in \mathcal{I}} \widehat{F}_M = 0)$. Consider the function:

$$(\widehat{F}_M)_{M \in \operatorname{ind} \operatorname{mod} \Lambda} : X \to (\mathbb{C}^*)^{|\operatorname{ind} \operatorname{mod} \Lambda|}.$$

The T action just rescales these values.

The locus where the $\widehat{F}_M$ are 1 is just the fibre over $(1, 1, \ldots, 1)$ of this map. Up to localization at the u_M, the locus where the $\widehat{F}_M = 1$ can be described as X/T. Since X is irreducible, so is the locus where the $\widehat{F}_M = 1$.

The A_1 example is already interesting. The $\widehat{F}$-polynomial is $u_{P(1)} + u_{P(1)[1]}$. (We recognize that setting this equal to 1 is equivalent to the u-equation for A_1, so that's reassuring.) The scaling is by a single $\mathbb{C}^*$, which rescales the two variables simultaneously, so also rescales the $\widehat{F}$-polynomial.

In the A_2 example, the action is by $(\mathbb{C}^*)^3$. One $\mathbb{C}^*$, parametrized by t, sends $u_{P(1)}$ to $tu_{P(1)}$, $u_{P(2)}$ to $t^{-1}u_{P(2)}$, and $u_{S(2)}$ to $tu_{S(2)}$, while not changing $u_{P(1)[1]}$ and $u_{P(2)[1]}$. We see that t sends $\widehat{F}_{S(1)}$ to $t\widehat{F}_{S(1)}$ while not changing the other two. We can similarly find two other $\mathbb{C}^*$ actions which act the same way with respect to the other two $\widehat{F}$ polynomials.

Since we know the $\overline{\mathcal{U}_{\geq 0}} \subseteq \mathcal{F} \subseteq \mathcal{U}$, the fact that $\mathcal{F}$ is irreducible means that $\mathcal{F} = \overline{\mathcal{U}_{\geq 0}}$. My impression is that the should also tell us that the above perspective should also tell us that the $\widehat{F}_M = 1$ equations are reduced (at least generically — the issue being that we localized by the u_M so as not to have to worry about what's happening when $u_M = 0$). If someone can help me out with how to state that argument precisely, I'd be grateful.

One thing that's perhaps interesting to notice is that the u-equations themselves are not homogeneous. So we can't use the same kind of "rescaling" argument directly on them.

3.6. **Functoriality of u-equations by quotients.** For this next section, it's necessary to think of our indexing set $\mathcal{I}$ as indecomposable 2-term complexes of projectives, rather than as modules and shifted projectives.

Let $A = \Lambda/I$. We write π for the functor $A \otimes_\Lambda -$ from Λ-modules to A-modules. Since π sends projective Λ-modules to projective A-modules, we can consider the action of π on $\mathcal{K}^2(\operatorname{proj} \Lambda)$.

I'm going to write $P(i)$ for the projective Λ-modules and $\overline{P}(i)$ for the corresponding projective A-module (and similarly for simple modules, etc.)

The following is non-surprising:

Lemma 3.9. $H^0(\pi(C)) = \pi(H^0(C))$

However, even if C is an indecomposable resolution, $\pi(C)$ can have $\overline{P}(i) \to 0$ summands!

Let's see this in action in type A_2. Let $\Lambda = \mathbb{C}[1 \leftarrow 2]$, and let A be the quotient of Λ by the class of the arrow.

Now consider the complex

$$C = P(1) \xrightarrow{i} P(2),$$

with i the inclusion. If we tensor by A, the image of the $P(1)$ is now zero, so

$$\pi(C) = \overline{P}(1) \xrightarrow{0} \overline{P}(2) = (\overline{P}(1) \to 0) \oplus (0 \to \overline{P}(2)).$$

Note that the cokernel of $\pi(C)$ is indeed $\overline{S}(2) = A \otimes S(2)$, as predicted by the lemma.

To summarize this example in the "wrong" language (i.e., using modules rather than complexes): $\pi(S(2)) = \overline{S}(2) \oplus \overline{P}(1)[1]$, where, if we are thinking in terms of modules, the second term is unexpected.

As usual, we have variables u_N for N in $\mathcal{I}$, the indecomposable complexes in $\mathcal{K}^2(\operatorname{proj}\Lambda)$ and also variables $\overline{u}_{\overline{M}}$ for $\overline{M}$ in $\overline{\mathcal{I}}$, the indecomposable object in $\mathcal{K}^2(\operatorname{proj} A)$.

Let $R = \mathbb{C}[u_N]/\langle \widehat{F}_M - 1\rangle$, and let $\overline{R} = \mathbb{C}[\overline{u}_{\overline{M}}]/\langle \widehat{F}_{\overline{M}} - 1\rangle$.

For $\overline{M} \in \overline{\mathcal{I}}$, and $\overline{N} \in \mathcal{K}^2(\operatorname{proj} A)$ write $m_{\overline{M}}(\overline{N})$ for the multiplicity of $\overline{M}$ as a direct summand of $\overline{N}$ Define:

$$\phi(\overline{u}_{\overline{M}}) = \prod_{N\in\mathcal{I}} u_N^{m_{\overline{M}}(\pi N)} \tag{1}$$

Theorem 3.10. *ϕ defines a ring map from $\overline{R}$ to R.*

That is to say, the products of u-variables for Λ given by (1) satisfy the equations $\widehat{F}_{\overline{M}} = 1$ for A.

Let's see this in action for the quotient from $1 \leftarrow 2$ which kills the arrow. As we already calculated, $\pi(P(1) \to P(2)) = (\overline{P}(1) \to 0) \oplus (0 \to \overline{P}(2)$. Otherwise $\pi(0 \to P(i)) = 0 \to \overline{P}(i)$, and similarly for $\pi(P(i) \to 0)$.

Thus, what we are saying is that the $A_1 \times A_1$ equations should be satisfied for the variables:

$$\begin{aligned}
u_{0\to P(1)} &= \frac{1}{1+y_1}\\
u_{0\to P(2)}u_{P(1)\to P(2)} &= \frac{1+y_1}{1+y_1+y_1y_2}\cdot\frac{1+y_1+y_1y_2}{(1+y_1)(1+y_2)} = \frac{1}{1+y_2}\\
u_{P(1)\to P(2)}u_{P(1)\to 0} &= \frac{1+y_1+y_1y_2}{(1+y_1)(1+y_2)}\cdot\frac{y_1(1+y_2)}{1+y_1+y_1y_2} = \frac{y_1}{1+y_1}\\
u_{0\to P(2)} &= \frac{y_2}{1+y_2}
\end{aligned}$$

We see that they are satisfied!

Proof. The proof is a calculation, showing that $\widehat{F}_{\overline{M}}$, after the desired substitution, equals $\widehat{F}$ for the same module $\overline{M}$, but viewed as a Λ module. So setting the $\widehat{F}$'s to 1 for Λ also sets the $\widehat{F}$'s for A to 1 (after the above substitution). □

Note that $[\pi N : \overline{M}]$ can be calculated solely on the basis of g_N if N is τ-rigid.

Proposition 3.11. *Let N be τ-rigid in $\mathcal{I}$. The expression for g_N as a positive combination of the generators of the minimal cone containing it gives the expression of πN as a sum of indecomposable objects from $\overline{\mathcal{I}}$.*

Proof. The key point to establish is that πN is again τ-rigid. Once we know that, πN is determined by its g-vector, which is the same as that of N, which proves the result. (This is not literally true if the quotient kills a projective summand entirely, but that doesn't cause a problem.) The fact that πN is again τ-rigid is a little calculation in the homotopy category. □

We can see that this holds in the example $Q = 1 \leftarrow 2$ which we have been investigating.

3.7. **Parenthesis: u-equations in general.** While it's very interesting to me that we get these u-equations from finite-dimensional algebras, it could also be interesting to think abstractly about the question: for what complete fans are there systems of u-equations for which the totally non-negative solutions recover the combinatorics of the fan? This is of practical interest to my physicist friends: encountering a particular fan, they would like to know if it has u-equations, whether or not they come from representation theory.

We know that $\Sigma(\Lambda)$ refines $\Sigma(A)$, i.e., every cone of $\Sigma(\Lambda)$ is contained in some cone of $\Sigma(A)$. It turns out that the picture of what happens when passing from the g-fan of Λ to the g-fan of A kills a single ray, sheds some light on how to create systems of u-equations with the expected dimensionality of solutions.

Suppose that Λ has exactly one more indecomposable representation than A, and suppose that $\Sigma(\Lambda)$ contains exactly one ray, ρ, not appearing in $\Sigma(A)$. Let σ be the smallest cone of $\Sigma(A)$ containing ρ. There is a well-understood picture for what the combinatorics here must look like: σ is replaced in $\Sigma(\Lambda)$ by the cones spanned by a facet of σ together with ρ; and any cone that contains σ is treated similarly.

Here are a couple of examples, where I have drawn an affine slice of part of the fan, rather than the whole fan (which in these cases would be three-dimensional).

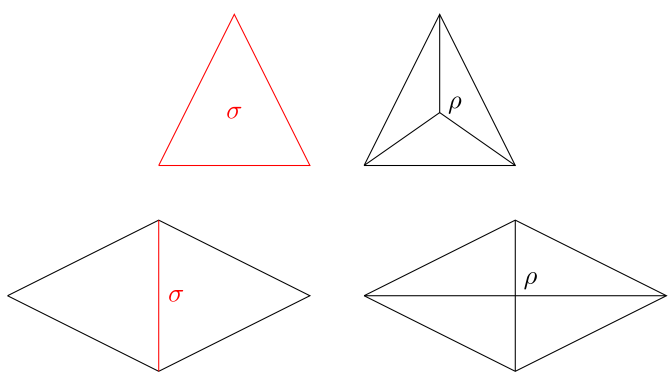

If σ were of dimension greater than 2, then the rays of σ are still in 2-dimensional cones with each other for Λ (i.e., after the blowup), which means that they are still all mutually compatible. But this would imply that for Λ, the generic presentation in the relative interior of σ would still be given as a sum of the presentations corresponding to the rays of σ, which contradicts the existence of a new presentation ρ. So the dimension of σ must be exactly 2.

Let the generators of σ are ξ_1 and ξ_2, corresponding to N_1 and N_2 for Λ, and to $\overline{N}_1$ and $\overline{N}_2$ for A. Since the generators of a cone of the g-fan form a basis for the lattice obtained by restricting the Grothendieck group to their linear span, it turns out that we must have $\rho = \xi_1 + \xi_2$.

This tells us that the $\overline{u}$-equations for A are satisfied by: $\overline{u}_{\overline{K}} = u_K$ for most $K \in \mathcal{I}$, together with $\overline{u}_{\overline{N}_1} = u_{N_1} u_M$ and $\overline{u}_{\overline{N}_2} = u_{N_2} u_M$. After a bit of work, it winds up that, to get the compatibility degrees for Λ from those of A, you just decree that $c(M, X) = \overline{c}(\overline{N}_1, X) + \overline{c}(\overline{N}_2, X)$, and all the other compatibilities are the same for K and $\overline{K}$, except that $c(N_1, N_2) = 1$, while $\overline{c}(\overline{N}_1, \overline{N}_2) = 0$ since they are compatible.

What this all suggests is the following: given a collection of u-equations corresponding to a fan in $\mathbb{R}^n$, with an n-dimensional family of solutions, could one create a new family of u-equations with an n-dimensional family of solutions by blowing up one two-dimensional face? And it turns out that the answer is yes, and the above discussion of torsion classes tell you exactly how to do it: given a 2-dimensional cone σ joining two rays ξ_1 and ξ_2, we blow it up, introducing a new ray ρ as above. If $\overline{u}_{\xi_1}$ and $\overline{u}_{\xi_2}$ are the u-variables downstairs associated to ξ_1, ξ_2, the corresponding variables upstairs, and the new variable, are:

$$\begin{aligned} u_\rho &= \overline{u}_{\overline{\xi}_1} + \overline{u}_{\overline{\xi}_1} - \overline{u}_{\overline{\xi}_1}\overline{u}_{\overline{\xi}_1} \\ u_{\xi_1} &= \overline{u}_{\overline{\xi}_1}/u_\rho \\ u_{\xi_2} &= \overline{u}_{\overline{\xi}_2}/u_\rho \end{aligned}$$

while all the other variables stay the same.

In particular, you can create the Koba–Nielsen u-equations this way, starting from the equations for $A_1 \times \cdots \times A_1$. (This is because there is a sequence of quotients, via Nakayama algebras, which remove one ray at a time from the g-fan of linearly oriented A_n to get down to $A_1 \times \cdots \times A_1$.) Is it possible that you can get the u-equations for any finite representation type algebra over an algebraically closed ground field, and such that all modules are τ-rigid, in this way? (I don't know if it's possible to get all the type A_n associahedra like this.)

In fact, there is a slightly more general thing one can do, inspired by the procedure above, which introduces a new u-varible which does not correspond to a boundary: pick two rays ξ_1 and ξ_2 in the downstairs fan which *don't* span a 2-dimensional cone. Again define the new variable ρ by $c(\rho, \mu) = \overline{c}(\overline{\xi}_1, \overline{\mu}) + \overline{c}(\overline{\xi}_2, \overline{\mu})$, and $c(\rho, \rho) = 2\overline{c}(\overline{\xi}_1, \overline{\xi}_2)$, and set $c(\xi_1, \xi_2) = \overline{c}(\overline{\xi}_1, \overline{\xi}_2) + 1$, with the other compatibilities as before. It turns out that, again, you get a set of u-equations defining an n-dimensional variety. Is this enough to get the u-equations for any finite-dimensional algebra of finite representation type?

3.8. **The g-fan toric variety and the u-equations.** Possible references for toric varieties are [CLS, Fu]. Here again, we are inspired by the approach of [AHL].

Let Σ be a complete unimodular fan in $\mathbb{R}^n$: that is to say, for any cone of Σ, the generators of the rays of σ span $\mathbb{R}\sigma \cap \mathbb{Z}^n$. (Weaker hypotheses on Σ are possible.)

Let P be a lattice polytope having Σ as its outer normal fan. Let $P \cap \mathbb{Z}^n = \{p_1, \ldots, p_s\}$. Then the toric variety X_Σ associated to Σ is the Zariski closure of the image of the map sending $T = (\mathbb{C}^*)^n$ to $\mathbb{P}^{s-1}$, sending t to $(t^{p_1}, \ldots, t^{p_s})$. I realize this is a bit brief as a definition. The feature is that the toric variety associated to Σ has an action of $(\mathbb{C}^*)^n$, with one orbit for each cone of Σ; a cone σ of dimension d corresponds to $(\mathbb{C}^*)^{n-d}$. The origin corresponds to the copy of $(\mathbb{C}^*)^n$ which we embedded into projective space to begin with.

X_Σ has a totally positive part: the closure of the image of $(\mathbb{R}_{>0})^n$ under the embedding into projective space. It is homeomorphic to P.

Write $F_\Lambda = \prod_{M \in \operatorname{ind} \Lambda} F_M$. Let P_Λ be the Newton polytope of F_Λ. Its outer normal fan is $\Sigma(\Lambda)$ by results of [F2, AH+]. Now use P_Λ to define $X_\Lambda = X_{\Sigma(\Lambda)}$.

Theorem 3.12. *$\mathcal{F}_\Lambda$ is isomorphic to the affine open subset of $X_{\Sigma(\Lambda)}$ where F_Λ is non-vanishing (where we interpret F_Λ as a linear combination of the coordinates on* *projective space, which corresponded to the lattice points in P_Λ, the Newton polytope of F_Λ.*

If Σ refines Σ', then there is a map of toric varieties $X_\Sigma \to X_{\Sigma'}$.

Theorem 3.13. *Let $A = \Lambda/I$. It is known that the g-vector fan of A coarsens that of Λ, and therefore there is an induced map $X_\Sigma \to X_A$. This map agrees with the map from $\mathcal{F}_\Lambda$ to $\mathcal{F}_A$ from the previous section, i.e., the following diagram commutes:*

$$\begin{tikzcd}
\mathcal{F}_\Lambda \arrow[r, hook] \arrow[d, two heads] & X_\Lambda \arrow[d, two heads] \\
\mathcal{F}_A \arrow[r, hook] & X_A
\end{tikzcd}$$

Corollary 3.14. *The totally positive parts of $\widetilde{\mathcal{M}}_\Lambda$ and X_Λ coincide.*

$\widetilde{\mathcal{M}}_\Lambda$ has a natural top form on it, namely the top form of the big torus contained in the toric variety X_Λ, which is $dy_1/y_1 \wedge \cdots \wedge dy_n/y_n$. It is natural to ask whether $\widetilde{\mathcal{M}}_\Lambda$ equipped with this differential form is a positive geometry in the sense of [ABL]. For a rather trivial reason, it isn't: namely, it is not projective. However, except for that issue, it is indeed a positive geometry, because its totally positive part agrees with the totally positive part of its ambient toric variety, and toric varieties are shown in [ABL] to be totally positive. (We are also implicitly using the fact that the boundaries of the totally positive part are again varieties of the same form.)

Acknowledgements. HT was partially supported by NSERC Discovery Grant RGPIN-2022-03960 and the Canada Research Chairs program, grant number CRC-2021-00120. Thanks to the organizers of the Autumn School on New developments in representation theory of algebras held at OIST in November 2024 for the opportunity to present this material.

References

[AIR] T. Adachi, O. Iyama, and I. Reiten, τ-tilting theory. Compositio Math. 150 (2014) no. 3, 415–452.

[AH+] Toshitaka Aoki, Akihiro Higashitani, Osamu Iyama, Ryoichi Kase, and Yuya Mizuno. Fans and polytopes in tilting theory I: Foundations. arXiv:2203.15213.

[ABL] Nima Arkani-Hamed, Yuntao Bai, and Thomas Lam. Positive geometries and canonical forms. *J. High Energy Phys.*, (11):039, front matter+121, 2017.

[AHL] N. Arkani-Hamed, S. He, T. Lam, Cluster configuration spaces of finite type. SIGMA 17 (2021) Paper no. 092.

[ABCP] I. Assem, T. Brüstle, G. Charbonneau-Jodoin, P.-G. Plamondon, Gentle algebras arising from surface triangulations. Algebra and Number Theory 4 (2010), no. 2, 201–229.

[ASS] I. Assem, D. Simson, A. Skowroński. Elements of the representation theory of associative algebras, Vol 1. Cambridge University Press, Cambridge, 2006.

[B] M. Brion, Representations of quivers. https://hal.science/cel-00440026/.

[BD+] T. Brüstle, G. Douville, K. Mousavand, H. Thomas, and E. Yıldırım, On the combinatorics of gentle algebras. Canadian J. Math. 72 (2020) no. 6, 1551–1580.

[CLS] David A. Cox, John B. Little, and Henry K. Schenck. *Toric varieties*, volume 124 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2011.

[DK] R. Dehy and B. Keller, On the combinatorics of rigid objects in 2-Calabi–Yau categories. IMRN 2008, no. 11, Art. ID rnn029.

[DF] H. Derksen and J. Fei, General presentations of algebras. Adv. Math. 278 (2015), 210–237.

[DWZ] H. Derksen, J. Weyman, and A. Zelevinsky, Quivers with potentials and their representations II: applications to cluster algebras. Journal of the AMS 23 (2010) no. 3, 749–790.

[DG] S. Domínguez and Ch. Geiss, A Caldero–Chapoton formula for generalized cluster categories. J. Alg. 399 (2014), 887–893.

[F1] J. Fei, Combinatorics of F-polynomials. IMRN 2023, no. 9, 7578–7615.

[F2] ———, Tropical F-polynomials and general presentations. J. LMS 107 (2023), no. 6, 2079–2120.

[FT] S. Fomin and D. Thurston, Cluster algebras and triangulated surfaces Part II: Lambda lengths. Mem. AMS 255 (2018) no. 1223.

[Fu] W. Fulton, Introduction to toric varieties. Princeton University Press, Princeton, 1993.

[J] G. Jasso, Reduction of τ-tilting modules and torsion pairs. IMRN 2015, no. 16, 7190–7237.

[K] V. Kac, Infinite root systems, representations of graphs, and invariant theory. Inventiones math. 56 (1980) 57–92.

[KN1] Z. Koba and H. B. Nielsen, Reaction amplitude for n-mesons: a generalizatino of the Veneziano–Bardakçi–Ruegg–Virasoro model. Nuclear Physics B (1969) no. 10, 633–655.

[KN2] ———, Manifestly crossing-invariant parametrization of n-meson amplitude. Nuclear Physics B (1969) no. 12, 517–536.

[KN3] ———, Generalized Veneziano model from the point of view of manifestly crossing-invariant parametrization. Zeitschrift für Physik 229 (1969), 243–264.