and restricting the cochain complex to the interior faces yields the minimal free resolution

$$0 \longleftarrow S^7 \longleftarrow S^9 \longleftarrow S^3 \longleftarrow 0$$

corresponding to 3 vertices, 9 edges, and 7 facets. □

Remark 6.24 The coScarf resolution is rarely the hull resolution of a cogeneric monomial ideal. It can happen, though, that they coincide, as they do for the permutohedron ideal. □

The coScarf resolution, though it is defined here to be "relative cocellular", is actually cellular, as well (like $\mathbb{K}^{\bullet}$ and $\mathbb{K}_{\bullet}$), as is evident for the permutohedron. To get the general idea, superimpose the hexagon (cellular version of the coScarf resolution for cogeneric permutohedron ideal) onto the barycentric subdivision of the triangle (Scarf resolution of generic tree ideal).

Just as the Scarf resolution is a special case of the hull resolution, the coScarf resolution is a special case of the *cohull resolution* [23], which does the same thing to the hull resolution of an arbitrary artinian monomial ideal as the coScarf resolution does to the Scarf resolution when the ideal is generic. In the cohull construction, the hull complex of the Alexander dual ideal $M^{\vee}$ is defined by its facets, which are labeled by the generators of M, while the vertices correspond to the irreducible components of M. This contrasts with the usual case, where the facets of the hull complex of M represent the irreducible components of M, while the vertices represent the minimal generators. Thus we complete our tour of Alexander duality, by showing how the statement about minimal generators vs. irreducible components and vertices vs. bounding half-spaces is really a deep truth rather than a convenient analogy: it is the principle underlying the interaction of Alexander duality with the geometric constructions of cellular resolutions.

7 Lecture VII: Monomial modules to lattice ideals

In this lecture and the next we study certain kinds of limits of monomial ideals. Here the limit will have an action of a *lattice* L—that is, a finitely generated subgroup of $\mathbb{Z}^n$. This establishes a relation between monomial ideals and binomial ideals: the quotient by L is a *lattice ideal*, of which a *toric ideal* is a special case. The cellular resolutions and genericity conditions of previous lectures will also prove to be of use. The general theory in this lecture was developed by D. Bayer and B. Sturmfels in [4], although genericity for lattice ideals is due to Sturmfels in collaboration with I. Peeva, and S. Popescu joined Bayer and Sturmfels for the work on Lawrence ideals [3].

7.1 Monomial modules

The Laurent polynomial ring $T = S[x_1^{-1}, \ldots, x_n^{-1}]$ is a module over S. An S-submodule M of T which is generated by Laurent monomials, $\mathbf{x}^{\mathbf{a}}$ with $\mathbf{a} \in \mathbb{Z}^n$, is called a *monomial module.* In general, such modules need not be minimally generated, but every monomial module M in this lecture will have a minimal generating set of monomials. If this set is finite, then M is a $\mathbb{Z}^n$-graded shift of a monomial ideal of S.

We can still draw pictures, but the usual staircase diagrams for monomial ideals become infinite staircases for monomial modules.

Example 7.1 Consider the monomial module in $k[x,y][x^{-1},y^{-1}]$ generated by the Laurent monomials $(\frac{x}{y})^i$ for $i \in \mathbb{Z}$. The staircase diagram really is a staircase, but an infinite one:

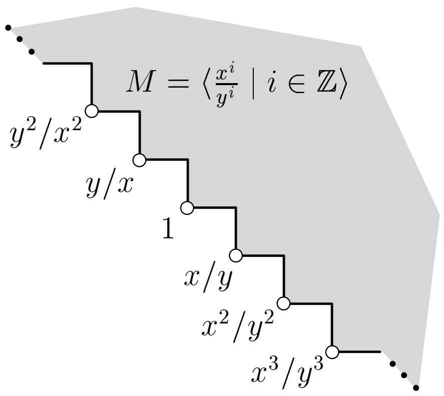

This monomial module can be thought of as the "mother of all powers of the maximal ideal": intersecting it with any shift of S produces a power of the maximal ideal. More generally, a monomial module can be thought of as the limit of the monomial ideals obtained by intersecting it with shifted positive orthants. □

The construction of the hull complex in Section 5.4 works here mutatis mutandis: fix $t \gg 0$, and form the unbounded n-dimensional polyhedron $P_t = \operatorname{conv}\{t^{\mathbf{a}} \mid \mathbf{x}^{\mathbf{a}} \in M\}$. The combinatorial type is independent of the large real number t, and its vertices are precisely the minimal generators of M (this is why we are assuming M has minimal generators). Again, a complex of free modules $\mathbb{F}_{\mathrm{hull}(M)}$ is defined, and the main result still holds; in fact, this is the generality in which it was originally published [4].

Theorem 7.2 *$\mathbb{F}_{\mathrm{hull}(M)}$ is a free resolution of M.*

Example 7.3 The hull complex hull(M) for the monomial module in the previous example is the real line with a vertex at each integer point. □

Example 7.4 A cell complex is said to be *locally finite* if every face meets finitely many others. In general, the hull complex of a monomial module need not be locally finite. For example, consider the monomial module M in three variables x, y, z:

$$M = \langle \tfrac{y}{x} \rangle + \langle (\tfrac{y}{z})^i \mid i \in \mathbb{Z} \rangle \,.$$

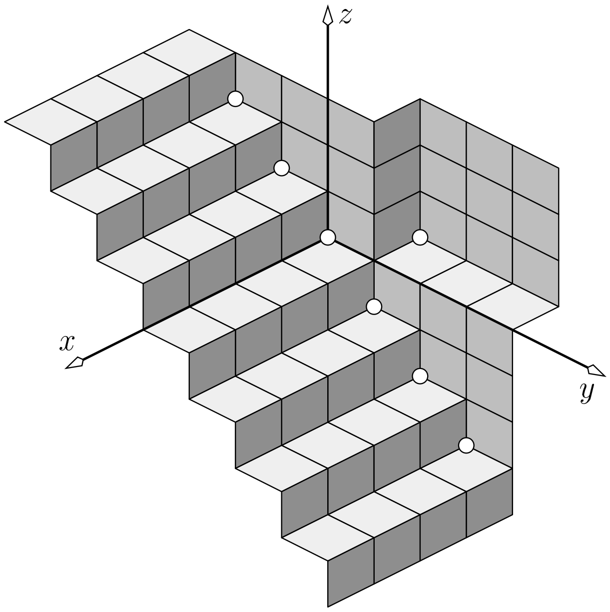

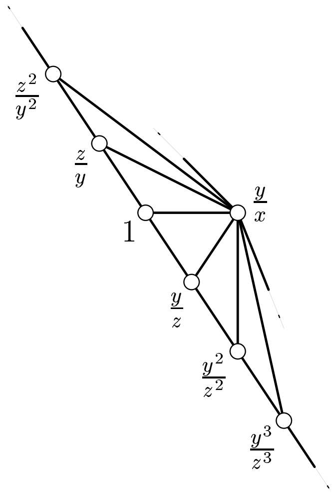

The vertex $\frac{y}{x}$ lies on infinitely many triangles of hull(M). $\square$

7.2 Lattice modules

In the all of the lectures before this one, we have been interested in monomial modules whose generating sets are finite. Here, we consider another interesting case, when the set of generating Laurent monomials forms a group under multiplication. To be precise, let $L \subset \mathbb{Z}^n$ be a sublattice whose intersection with $\mathbb{N}^n$ is $\{\mathbf{0}\}$. This ensures the existence of a functional with strictly positive coordinates that vanishes on L.

Definition 7.5 The *lattice module* M_L is the S-submodule of $T = S[x_1^{-1}, \ldots, x_n^{-1}]$ generated by $\{\mathbf{x}^{\mathbf{a}} \mid \mathbf{a} \in L\}$.

The hypothesis on L guarantees that the elements of L form a *minimal* generating set for M_L.

Example 7.6 The monomial module in Example 7.1 is actually the lattice module M_L for the lattice $\ker(1,1)$. More generally, the "mother of all maximal ideals" in n variables is the lattice module for $L = \ker(1,\ldots,1)$. For a picture of a finite part of the staircase of $\ker(1,1,1)$, where $n = 3$, see Example 5.10. $\square$

Because the generating monomials form a group L, that group acts on the monomial module M_L, the action being given by $\mathbf{a} \cdot \mathbf{x}^{\mathbf{b}} = \mathbf{x}^{\mathbf{a}+\mathbf{b}}$ for $\mathbf{a} \in L$ and $\mathbf{x}^{\mathbf{b}} \in M_L$. What would be really nice is if we could find a whole free resolution of M_L which is acted on by L. Such an *equivariant free resolution* is provided by the hull resolution. The point is that the lattice L permutes the faces of hull(M).

Example 7.7 The hull complex of the "mother of all maximal ideals" in two variables (Example 7.3) is obviously acted on by $L \cong \mathbb{Z}$. In three variables, the hull complex is the tesselation of $\mathbb{R}^2$ by triangles, a part of which is depicted in Example 5.10. Note that in the case of three variables, the action of $L = \ker(1,1,1)$ on the hull complex is easiest to see in the picture if we view the tesselation as actually sitting in $\mathbb{R} \otimes L \subset \mathbb{R}^3$ instead of in the plane. □

Calculating the hull complex is actually a finite thing to do (although it is an open problem to determine an efficient algorithm), even though it has infinitely many cells. This is because of the following minor miracle:

Proposition 7.8 *The hull complex of a lattice module is locally finite.*

In fact, one produces an explicit finite set of edges (a *Graver basis* for L) such that the vertex $\mathbf{0} \in L$ meets at most these edges, and hence $\mathbf{0}$ meets finitely many faces altogether. But every vertex of $\mathrm{hull}(M_L)$ is equivalent to $\mathbf{0}$ under the action of L, so we only need the local behavior of $\mathrm{hull}(M_L)$ around $\mathbf{0}$ to determine the entire hull complex. Thus we reduce the calculation of $\mathrm{hull}(M_L)$ to the computation of $\mathrm{hull}(I)$ for a monomial module I generated by a finite set of elements of L. In other words, the hull complex of a monomial ideal is enough information to determine $\mathrm{hull}(M_L)$ *modulo the action of L on the faces*, which has only finitely many orbits.

7.3 Genericity

A monomial module $M \subset T$ is *generic* if all its minimal first syzygies $\mathbf{x}^{\mathbf{a}}\mathbf{e}_i - \mathbf{x}^{\mathbf{b}}\mathbf{e}_j$ have full support, i.e. every variable x_ℓ appears either in $\mathbf{x}^{\mathbf{a}}$ or $\mathbf{x}^{\mathbf{b}}$. This definition is the essence behind genericity for monomial ideals, although for ideals there are "boundary effects" coming from the fact that $\mathbb{N}^n$ is a special subset of $\mathbb{Z}^n$. To be precise, the genericity condition on the minimal first syzygies $\mathbf{x}^{\mathbf{a}}\mathbf{e}_i - \mathbf{x}^{\mathbf{b}}\mathbf{e}_j$ of an ideal requires only that $\mathrm{supp}(\mathbf{x}^{\mathbf{a}+\mathbf{b}}) = \mathrm{supp}(\mathrm{lcm}\,(m_i, m_j))$, as opposed to $\mathrm{supp}(\mathbf{x}^{\mathbf{a}+\mathbf{b}}) = \{1, \ldots, n\}$ for monomial modules. This definition allows us to treat the boundary exponent 0 differently than the strictly positive exponents coming from the interior of $\mathbb{N}^n$. Just like the hull complex, the Scarf complex defined earlier for monomial ideals makes sense for monomial modules, too, as does the theorem on free resolutions of generic objects.

Theorem 7.9 *For a generic monomial module M, the following coincide:*

1. *the Scarf complex of M;*
2. *the hull resolution of M; and*
3. *the minimal free resolution of M.*

Example 7.10 The lattice $L = \ker([4\ 3\ 5]) \subset \mathbb{Z}^3$ yields a *generic lattice module*

$$M_L = \langle x^i y^j z^k \mid 4i + 3j + 5k = 0 \rangle$$

in three variables whose staircase is pictured below.

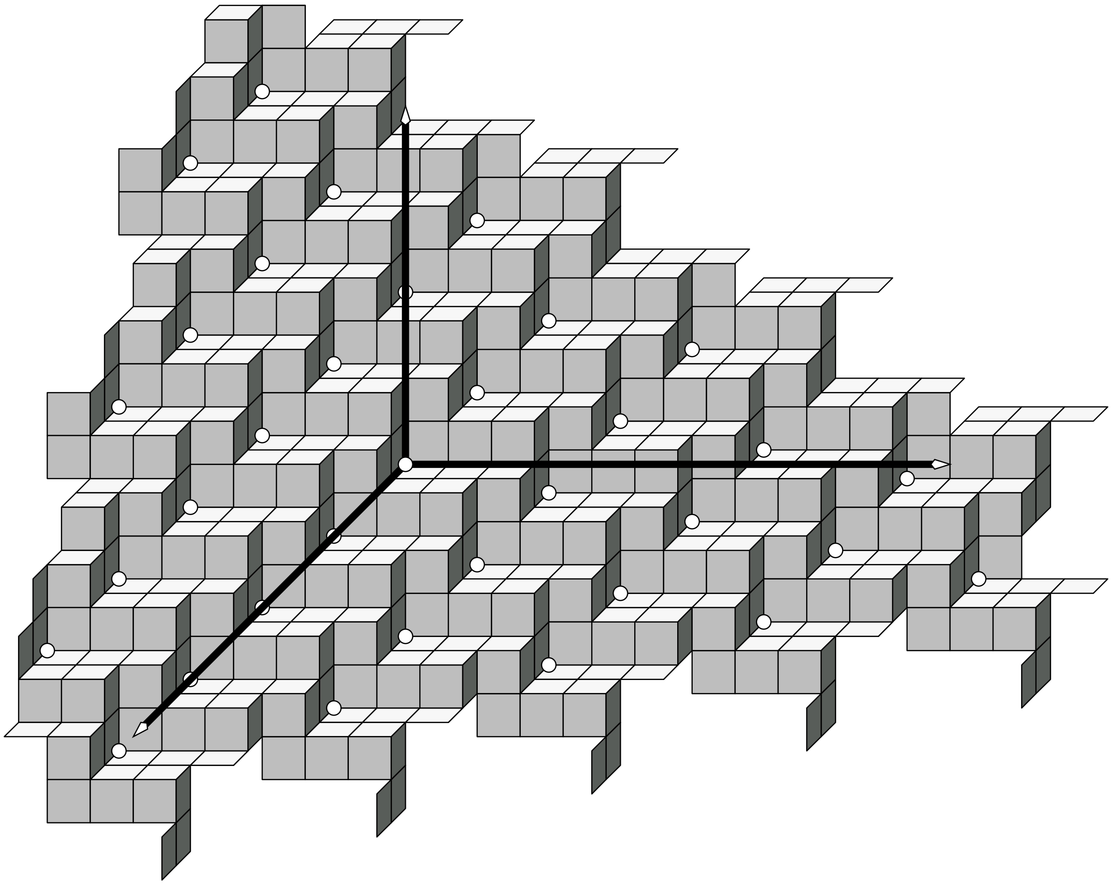

The hull = Scarf complex is a triangulation of $\mathbb{R} \otimes L$ using L for vertices:

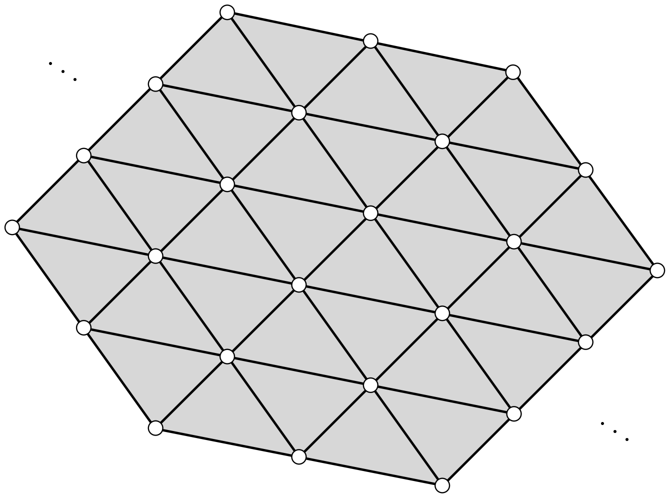

The labeling on every pair of up and down triangles is obtained from the representative labeling

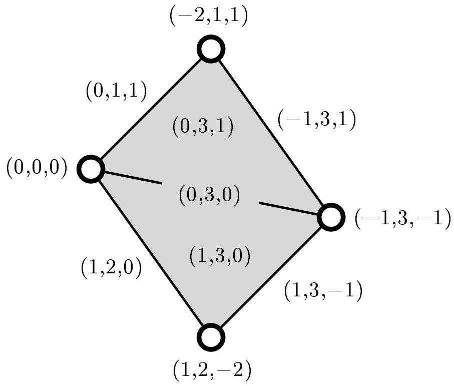

by adding some vector in L to all of the labels. $\square$

7.4 Lattice ideals

For a lattice $L \subset \mathbb{Z}^n$, define the *lattice ideal*

$$I_L = \langle \mathbf{x}^{\mathbf{a}} - \mathbf{x}^{\mathbf{b}} \mid \mathbf{a} - \mathbf{b} \in L \text{ and } \mathbf{a}, \mathbf{b} \in \mathbb{N}^n \rangle \subset S.$$

Assume, from now on, that $L \cap \mathbb{N}^n = \{\mathbf{0}\}$, so that the ideal is homogeneous with respect to some grading where $\deg(x_i)$ is a positive integer (use a positive functional which vanishes on L to define the $\mathbb{Z}$-grading). Such ideals arise in a wide variety of contexts, from integer programming and toric geometry to mirror symmetry and hypergeometric differential equations (see [33] and [31]). Given L, the fundamental things we would like to know about I_L are:

1. generators for I_L;
2. the $\mathbb{Z}^n/L$*-graded* Hilbert series of S/I_L, as a rational function; and
3. a (minimal) free resolution of S/I_L over S.

Of course, $3 \Rightarrow 2 \Rightarrow 1$, so we'll aim for free resolutions.

The essential idea is to express the lattice ideal I_L as a quotient of the monomial module M_L by the action of L. In order to do that, let's formalize the action by introducing the *group algebra* $S[L]$ of L over S. Explicitly, this is the algebra

$$S[L] = k[\mathbf{x}^{\mathbf{a}}\mathbf{z}^{\mathbf{b}} \mid \mathbf{a} \in \mathbb{N}^n \text{ and } \mathbf{b} \in L] \subset S[z_1^{\pm 1}, \ldots, z_n^{\pm 1}],$$

which carries a $\mathbb{Z}^n$-grading via $\deg(\mathbf{x}^{\mathbf{a}}\mathbf{z}^{\mathbf{b}}) = \mathbf{a} + \mathbf{b}$. The action of L on the lattice module $M_L = \langle \mathbf{x}^{\mathbf{a}} \mid \mathbf{a} \in L \rangle \subset T$ defined in the previous lecture can be expressed as $\mathbf{b} \cdot \mathbf{x}^{\mathbf{a}} = \mathbf{x}^{\mathbf{a}+\mathbf{b}}$ for $\mathbf{b} \in L$ and $\mathbf{x}^{\mathbf{a}} \in M_L$. This action is reformulated using the group algebra by stipulating that $\mathbf{x}^{\mathbf{a}}\mathbf{z}^{\mathbf{b}} = \mathbf{x}^{\mathbf{a}+\mathbf{b}}$, making M_L into a $\mathbb{Z}^n$-graded module

$$M_L \cong S[L]/\langle \mathbf{x}^{\mathbf{a}} - \mathbf{x}^{\mathbf{b}}\mathbf{z}^{\mathbf{a}-\mathbf{b}} \mid \mathbf{a} - \mathbf{b} \in L \text{ and } \mathbf{a}, \mathbf{b} \in \mathbb{N}^n \rangle$$

over $S[L]$, with the image of $\mathbf{x}^{\mathbf{a}}$ having $\mathbb{Z}^n$-graded degree $\mathbf{a}$, as usual. In fact, any $\mathbb{Z}^n$-graded S-module M that has an equivariant action of L (i.e. commuting with the action of S) such that $\mathbf{b} \in L$ acts as a homomorphism of degree $\mathbf{b}$ is naturally a module over $S[L]$. This is really quite a general statement. To be precise, the category of L-equivariant $\mathbb{Z}^n$-graded S-modules is isomorphic to the category

$$\mathcal{A} = \{\mathbb{Z}^n\text{-graded } S[L]\text{-modules}\}\,.$$

How do we define the quotient of an L-equivariant module = $\mathbb{Z}^n$-graded $S[L]$-module M by the action of L? We would like to identify $m \in M$ with $\mathbf{z}^{\mathbf{b}} \cdot m$ whenever $\mathbf{b} \in L$, so that the quotient is an S-module whose elements are orbits of the action of L on M_L. When $M = S[L]$ itself, this quotient is

$$\begin{aligned} S[L]/L &= S[L]/\langle \mathbf{x}^{\mathbf{a}}\mathbf{z}^{\mathbf{b}} - \mathbf{x}^{\mathbf{a}} \mid \mathbf{a} \in \mathbb{N}^n \text{ and } \mathbf{b} \in L\rangle \\ &= S[L]/\langle \mathbf{z}^{\mathbf{b}} - 1 \mid \mathbf{b} \in L\rangle \\ &\cong S\,. \end{aligned}$$

But this copy of S is no longer $\mathbb{Z}^n$-graded, because $\mathbf{x}^{\mathbf{a}}$ and $\mathbf{x}^{\mathbf{a}}\mathbf{z}^{\mathbf{b}}$, which have different $\mathbb{Z}^n$-graded degrees $\mathbf{a}$ and $\mathbf{a}+\mathbf{b}$, map to the same element $\mathbf{x}^{\mathbf{a}}$. On the other hand, all of the preimages in $S[L]$ of $\mathbf{x}^{\mathbf{a}} \in S$ have $\mathbb{Z}^n$-graded degrees which are congruent modulo L, so this copy of S is $\mathbb{Z}^n/L$-*graded*, with $\mathbf{x}^{\mathbf{a}}$ having $\mathbb{Z}^n/L$-graded degree $\mathbf{a}$ (mod L).

For more general $\mathbb{Z}^n$-graded $S[L]$-modules M, our quotient module "M/L" will similarly be obtained by "setting all elements $\mathbf{z}^{\mathbf{b}}$ for $\mathbf{b} \in L$ equal to 1". Algebraically, this is just tensoring M over $S[L]$ with $S = S[L]/\langle \mathbf{z}^{\mathbf{b}} - 1 \mid \mathbf{b} \in L\rangle$,

$$M/L = M \otimes_{S[L]} S[L]/L = M \otimes_{S[L]} S\,.$$

As with $S[L]/L$, the quotient M/L is no longer $\mathbb{Z}^n$-graded, but only $\mathbb{Z}^n/L$-graded. This tensor product therefore defines a functor

$$\pi : \mathcal{A} \longrightarrow \mathcal{B} = \{\mathbb{Z}^n/L\text{-graded } S\text{-modules}\}$$

of categories. Our motivating example is M_L, whose image in $\mathcal{B}$ is obtained from its presentation as a quotient of $S[L]$ given above by setting all occurrences of $\mathbf{z}$ to 1:

$$\pi(M_L) = S[L]/\langle \mathbf{x}^{\mathbf{a}} - \mathbf{x}^{\mathbf{b}}\mathbf{z}^{\mathbf{a}-\mathbf{b}} \mid \mathbf{a} - \mathbf{b} \in L \text{ and } \mathbf{a}, \mathbf{b} \in \mathbb{N}^n\rangle|_{\mathbf{z}=1} = S/I_L\,.$$

Therefore we have achieved our goal of writing S/I_L as the quotient of M_L by the action of L. The great thing about π is that it doesn't forget anything significant.

Theorem 7.11 *The functor $\pi : \mathcal{A} \to \mathcal{B}$ sending $M \mapsto M/L$ is an equivalence of categories.*

The main idea of the proof is that if L acts equivariantly on a $\mathbb{Z}^n$-graded S-module in such a way that $\mathbf{b} \in L$ acts as a homomorphism of degree $\mathbf{b}$, then the action of L is necessarily free—no nonzero $\mathbf{b} \in L$ can have any nonzero fixed points. Therefore, an "inverse" to π is given by taking the universal cover.

Now we can use the functoriality of π to do for free resolutions what it did for modules.

Corollary 7.12 *If $\mathbb{F}$ is a free resolution of M_L over $S[L]$, then $\pi(\mathbb{F})$ is a $\mathbb{Z}^n/L$-graded resolution of S/I_L over S.*

And what is a resolution of M_L over $S[L]$? It's just a resolution of M_L as an S-module along with a free action of L. These exist because π is an equivalence. But we have already noted that L acts on nice resolutions of M_L such as the hull resolution.

Definition 7.13 The *hull resolution* of S/I_L is $\pi(\mathbb{F}_{\mathrm{hull}(M_L)})$.

Theorem 7.14 *The hull resolution of S/I_L is a finite free resolution, of length $\leq n$.*

Proof: L acts freely on $\mathrm{hull}(M_L)$, which implies that $\mathbb{F}_{\mathrm{hull}(M)}$ is a free $S[L]$-module. Since π(free $S[L]$-module) is a free S-module, the hull resolution of S/I_L is a resolution by free S-modules. The finiteness is because of the local finiteness of Proposition 7.8. And the length is $\leq n$ because $\mathrm{hull}(M)$ is a polyhedron inside $\mathbb{R}^n$. □

The rest of this section (and lecture) is devoted to examples.

Example 7.15 The hull resolution of a generic lattice ideal is π(Scarf complex). For instance, the ideal for the image of the affine monomial curve $t \mapsto (t^4, t^3, t^5)$ is the generic lattice ideal

$$I_L = \langle xy^2 - z^2, xz - y^3, yz - x^2 \rangle \,,$$

where L is the kernel of the matrix $[4\ 3\ 5]$ over $\mathbb{Z}$. The corresponding lattice module is the one whose staircase and Scarf complex are drawn in Example 7.10. Modulo the lattice, we see that the resolution of S/I_L is a torus

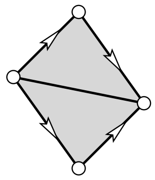

whose fundamental domain is labeled with vectors in Example 7.10. (So π really is a universal cover, as suggested by the proof of Theorem 7.11.) Under the functor π, the resolution of S/I_L is

$$0 \longleftarrow S \xleftarrow{\begin{pmatrix} xy^2-z^2 & xz-y^3 & yz-x^2 \end{pmatrix}} S^3 \xleftarrow{\begin{pmatrix} y & x \\ x & z \\ z & y^2 \end{pmatrix}} S^2 \longleftarrow 0\,.$$

To get the $\mathbb{Z}^3/L \cong \mathbb{Z}$-graded Hilbert series, we simply take the alternating sum of the $\mathbb{Z}^3/L$-degrees of the faces of $\mathrm{hull}(M_L)/L$ and divide by the appropriate denominator. Each $\mathbb{Z}^3/L$-degree is the dot product of the label in Example 7.10 with $(4, 3, 5)$:

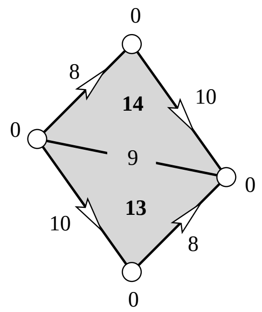

The graded Euler characteristic gives the numerator of the Hilbert series

$$\frac{1-t^8-t^9-t^{10}+t^{13}+t^{14}}{(1-t^4)(1-t^3)(1-t^5)} = \frac{1}{1-t} - t - t^2 = \sum\{t^i \mid i \in \mathbb{N}\{4,3,5\}\}$$

of S/I_L, where $\mathbb{N}\{4,3,5\}$ is the submonoid of $\mathbb{N}$ generated by $\{4,3,5\}$. The denominator comes from the $\mathbb{Z}^3/L$-graded Hilbert series $\frac{1}{(1-t^4)(1-t^3)(1-t^5)}$ of $k[x,y,z]$. □

Example 7.16 Things become much more complicated in four dimensions. The smallest codimension 1 generic lattice module in 4 variables is determined by the lattice $L = \ker([20\ 24\ 25\ 31]) \subset \mathbb{Z}^4$. The lattice ideal I_L is the ideal of the monomial curve $t \mapsto (t^{20}, t^{24}, t^{25}, t^{31})$ in affine four-space. The group algebra is $S[L] = k[a,b,c,d][\mathbf{z}^{\mathbf{a}} \mid \mathbf{a} \in L]$, and

$$\begin{aligned} M_L \quad = \quad & S[L]/\langle a^4 - bcd\,\mathbf{z}^*, a^3c^2 - b^2d^2\,\mathbf{z}^*, a^2b^3 - c^2d^2\,\mathbf{z}^*, ab^2c - d^3\,\mathbf{z}^*, \\ & b^4 - a^2cd\,\mathbf{z}^*, b^3c^2 - a^3d^2\,\mathbf{z}^*, c^3 - abd\,\mathbf{z}^*\rangle\,, \end{aligned}$$

where, for instance the * in $a^4 - bcd\,\mathbf{z}^*$ is the vector in L which is 4 times the first generator minus 1 times each of the second, third, and fourth generators. The hull = Scarf = minimal resolution of S/I_L is

$$0 \longleftarrow S \longleftarrow S^7 \longleftarrow S^{12} \longleftarrow S^6 \longleftarrow 0\,.$$

Up to the action of L, there are 6 tetrahedra corresponding to the second syzygies and 12 triangles corresponding to the first syzygies. □

Example 7.17 Suppose that L is unimodular, i.e. for all $\sigma \subseteq \{1,\ldots,n\}$ the coordinate projection $\mathbb{Z}^n/(L + \sum_{i\in\sigma}\mathbb{Z}\mathbf{e}_i)$ of $\mathbb{Z}^n/L$ is torsion free. Consider the Lawrence lifting, $\Lambda(L) = \{(\mathbf{u}, -\mathbf{u}) \in \mathbb{Z}^{2n} \mid \mathbf{u} \in L\}$, and its corresponding lattice ideal

$$I_{\Lambda(L)} = \langle \mathbf{x}^{\mathbf{a}}\mathbf{y}^{\mathbf{b}} - \mathbf{x}^{\mathbf{b}}\mathbf{y}^{\mathbf{a}} \mid \mathbf{a} - \mathbf{b} \in L\rangle \subset k[x_1,\ldots,x_n,y_1,\ldots,y_n]\,.$$

These *unimodular Lawrence ideals* have a remarkable rigidity property: all of their initial ideals are squarefree. In fact, this condition is equivalent to unimodularity for Lawrence ideals.

The hull resolution of $I_{\Lambda(L)}$ is not necessarily minimal, even if L is unimodular. However, the minimal resolution does come from a cellular resolution of $M_{\Lambda(L)}$, and is described by the following combinatorial construction.

Step 1. Take the infinite hyperplane arrangement $\{x_i = j \mid i = 1, \ldots, n,\ j \in \mathbb{Z}\}$.

Step 2. Let $\mathcal{H}_L$ be its intersection with $L \otimes \mathbb{R}$.

Step 3. Form the quotient $\mathcal{H}_L/L$.

In particular, the i-faces of $\mathcal{H}_L/L$ are in 1-1 correspondence with the minimal i-th syzygies of $I_{\Lambda(L)}$. A particular example of the resolution described here is the Eagon-Northcott complex for the 2×2 minors of a generic $2 \times n$ matrix [3]. □

8 Lecture VIII: Local cohomology

Among the first homological objects to be calculated explicitly for squarefree monomial ideals I were the local cohomology modules of S/I with support on $\mathfrak{m}$. Their Hilbert series were described in an unpublished theorem of M. Hochster (which finally found its way into [32]), and their module structure was later described explicitly by H. G. Gräbe [17]. On the other hand, local cohomology with support on the irrelevant monomial ideal in the Cox homogeneous coordinate ring can be used to compute sheaf cohomology on toric varieties, in analogy to the case of projective space [13, 28]. This application has prompted research into local cohomology with support on monomial ideals in general. N. Terai [38] made the calculation of the Hilbert series of the local cohomology modules of S with support on I, inspired by the techniques of Hochster; at the same time, M. Mustaţă [29] found the module structure. Then, using generalizations of Alexander duality for modules, E. Miller [24] showed that the answers provided by Terai and Mustaţă are actually equivalent to those provided by Hochster and Gräbe. Our aim in this final lecture is to present these results, along with the necessary background, in the spirit of geometry from previous lectures.

8.1 Preliminaries

The Čech complex. Local cohomology has many equivalent definitions, each with its own role in the theory. Hochster's local cohomology formula is proved using the most accessible of these, the *Čech complex* (more accurately called the *stable Koszul complex*),

$$\mathbb{C}^{\boldsymbol{\cdot}} : 0 \to S \to \bigoplus_{i=1}^{n} S[x_i^{-1}] \to \cdots \to \bigoplus_{|\sigma|=r} S[x^{-\sigma}] \to \cdots \to S[x_1^{-1}, \ldots, x_n^{-1}] \to 0$$

(see Section 0.1 for notation). This is to be considered as a cochain complex (upper indices increasing from the copy of S sitting in cohomological degree 0), with the map between the summands $S[\mathbf{x}^{-\sigma}] \to S[\mathbf{x}^{-(\sigma \cup i)}]$ in $\mathbb{C}^{\boldsymbol{\cdot}}$ being $\operatorname{sign}(i, \sigma \cup i)$ times the canonical localization homomorphism.

One way to think about the Čech complex in a $\mathbb{Z}^n$-graded way is to replace each summand $S[-\sigma]$ in the Koszul complex $\mathbb{K}^{\boldsymbol{\cdot}}$ by the localization $S[\mathbf{x}^{-\sigma}]$. That this replacement is possible can be seen from monomial matrices, whose scalar entries can be used just as easily as maps between localizations as between free modules; in fact