

Algorithms for graded injective resolutions and local cohomology over semigroup rings

David Helm and Ezra Miller

Mathematics Department, Harvard University, Cambridge, MA, USA

Mathematical Sciences Research Institute, Berkeley, CA, USA

(current address: School of Mathematics, University of Minnesota, Minneapolis, MN, USA)

(Received 22 March 2004)

Let Q be an affine semigroup generating \mathbb{Z}^d , and fix a finitely generated \mathbb{Z}^d -graded module M over the semigroup algebra $k[Q]$ for a field k . We provide an algorithm to compute a minimal \mathbb{Z}^d -graded injective resolution of M up to any desired cohomological degree. As an application, we derive an algorithm computing the local cohomology modules $H_I^i(M)$ supported on any monomial (that is, \mathbb{Z}^d -graded) ideal I . Since these local cohomology modules are neither finitely generated nor finitely cogenerated, part of this task is defining a finite data structure to encode them.

Contents

- 1 Introduction
 - 2 Effective irreducible hulls
 - 3 Computing with irreducible hulls
 - 4 Computing injective resolutions
 - 5 Sector partitions from injectives
 - 6 Computing sector partitions
 - 7 Computing local cohomology with monomial support
- References

1. Introduction

Injective resolutions are fundamental homological objects in commutative algebra. For general noetherian rings with arbitrary gradings, however, injective modules are so big, and injective resolutions so intractable, that effective computations are never made using them. But when the ring in question is an affine semigroup ring of dimension d , the natural grading by \mathbb{Z}^d is substantially better behaved: \mathbb{Z}^d -graded injective modules can be expressed polyhedrally and are therefore quite explicit. In this paper we provide algorithms to compute \mathbb{Z}^d -graded injective resolutions over affine semigroups rings. Part of this task is finding a finite data structure to express the output.

As an application, we provide an algorithm to compute the local cohomology, supported

on an arbitrary monomial ideal, of a finitely generated \mathbb{Z}^d -graded module over a normal affine semigroup ring. As far as we are aware, this is the first algorithm to compute local cohomology for any general class of modules over any class of nonregular rings.

Our motivation was to make a systematic study of conditions on a support ideal that cause the local cohomology of its ambient ring to have infinite Bass numbers. In particular, when does the local cohomology contain an infinite-dimensional vector subspace annihilated by a maximal ideal? A counterexample to Grothendieck’s conjectured answer of ‘never’ was provided by (Har70), and our motivation was to characterize when Grothendieck’s conjecture fails. That such infinite behavior occurs only in nonregular contexts suggested that we work over affine semigroup rings, which are among the simplest singular rings. We have not yet implemented the algorithms in this paper, as doing so would only be the first step to providing examples of the infinite-dimensional socle phenomenon: it still remains to find an algorithm computing the \mathbb{Z}^d -graded socle degrees.

To make our context precise, let $Q \subset \mathbb{Z}^d$ be an affine semigroup, that is, a finitely generated submonoid of \mathbb{Z}^d . We assume that Q is sharp, meaning that Q has no units, and that Q generates \mathbb{Z}^d as a group. Consider the semigroup algebra $k[Q] = \bigoplus_{a \in Q} k \cdot \mathbf{x}^a$ over a field k . The modules that concern us comprise the category \mathcal{M} of \mathbb{Z}^d -graded modules $H = \bigoplus_{\alpha \in \mathbb{Z}^d} H_\alpha$ for which there exists a bound independent of α on the dimensions of the graded pieces H_α as vector spaces over k . The injective objects in \mathcal{M} are described in Section 2, and every finitely generated \mathbb{Z}^d -graded module lies in \mathcal{M} . Our main theorem concerning injectives is the following.

THEOREM 1.1. *Fix a finitely generated \mathbb{Z}^d -graded module M over an affine semigroup ring $k[Q]$ and an integer $n \geq 0$. The first n stages in a minimal \mathbb{Z}^d -graded injective resolution of M can be expressed in a finite, algorithmically computable data structure.*

A more precise version, along with a pointer to the algorithms that do the job, is stated in Theorem 4.7. The data structure consists of a list of *monomial matrices*, as we define in Section 2, generalizing those for $Q = \mathbb{N}^d$ in (Mil00). The idea of the algorithm in Theorem 1.1 is to do all computations using *irreducible resolutions* (Mil02) as faithful approximations to injective resolutions. Background on *irreducible hulls* is presented in Section 2; the algorithms for working with them constitute Section 3. The derivation of an algorithm for injective resolutions is then completed in Section 4.

Even more seriously than is the case with injective resolutions, a substantial part of building an algorithm to compute local cohomology is finding a finite data structure to express the output. Indeed, unlike injectives in our category \mathcal{M} , and in stark contrast with the regular case (even without a grading (HS93, Lyu93)), the local cohomology $H_I^i(M)$ often has neither a finite generating set *nor* a *finite cogenerating set* (Har70, HM03). This remains true even when M is finitely generated and $I \subseteq k[Q]$ is a \mathbb{Z}^d -graded ideal—that is, generated by monomials. Our solution is to decompose \mathbb{Z}^d into tractable regions on which the local cohomology is constant.

DEFINITION 1.2. Suppose H is a \mathbb{Z}^d -graded module over an affine semigroup ring $k[Q]$. A **sector partition** of H is

1. a finite partition $\mathbb{Z}^d = \bigcup_{S \in \mathcal{S}} S$ of the lattice \mathbb{Z}^d into **sectors**, each of which is required to consist of the lattice points in a finite disjoint union of rational polyhedra defined as intersections of half-spaces for hyperplanes parallel to facets of Q ;

2. a finite-dimensional vector space H_S for each sector $S \in \mathcal{S}$, along with isomorphisms $H_\alpha \rightarrow H_S$ for all \mathbb{Z}^d -graded degrees $\alpha \in S$; and
3. vector space homomorphisms $H_S \xrightarrow{\mathbf{x}^{T-S}} H_T$ whenever there exist $\alpha \in S$ and $\beta \in T$ satisfying $\beta - \alpha \in Q$, such that for all choices of α and β , the diagram commutes:

$$\begin{array}{ccc} H_\alpha & \xrightarrow{\mathbf{x}^{\beta-\alpha}} & H_\beta \\ \downarrow & & \downarrow \\ H_S & \xrightarrow{\mathbf{x}^{T-S}} & H_T \end{array}$$

Write $\mathcal{S} \vdash H$ to indicate the above sector partition. (The commutativity of the above diagram implies immediately that \mathbf{x}^{S-S} is the identity, and that $\mathbf{x}^{R-T} \mathbf{x}^{T-S} = \mathbf{x}^{R-S}$.)

The finite data structure of a sector partition $\mathcal{S} \vdash H$, including the spaces H_S and the maps \mathbf{x}^{T-S} , clearly suffice to reconstruct H up to isomorphism. The second half of this paper is devoted to computing sector partitions for H when $H = H_I^i(M)$ is a local cohomology module.

THEOREM 1.3. *For any finitely generated \mathbb{Z}^d -graded module M over a normal semigroup ring $\mathbb{k}[Q]$ and any monomial ideal I , each local cohomology module $H_I^i(M)$ has an algorithmically computable sector partition $\mathcal{S} \vdash H_I^i(M)$.*

Section 5 demonstrates how sector partitions arise for the cohomology of any complex of injectives over a normal semigroup ring. Algorithms for producing these sector partitions, particularly the expressions of sectors as unions of polyhedral sets of lattice points, occupy Section 6. The proof of Theorem 1.3, by expressing local cohomology as the cohomology of a complex of injectives (algorithmically computed by Theorem 1.1) in the usual way, occurs in Section 7. That section also treats complexity issues. The main thrust is that for fixed dimension d , the running times of our algorithms are all polynomial in the Bass numbers of the finitely generated input module M and the number of facets of Q , times the usual factor arising from the complexity of Gröbner basis computation, where it occurs. If d is allowed to vary, then the numbers of polyhedra comprising sectors increase exponentially with d .

Theorem 1.3 allows the computation of many features of local cohomology modules. For example, Hilbert series simply record the vector space dimensions in each of the finitely many sectors $S \in \mathcal{S}$. Our algorithms can actually calculate these dimensions without computing the maps in part 3 of Definition 1.2, making it easier to determine when (for example) $H_I^i(M)$ is nonzero. Future algorithmic methods (currently open problems) include the calculation of associated primes and locations of socle degrees (even if there are infinitely many) using a sector partition as input. In particular, because of the finiteness of the number of polyhedra partitioning sectors, we believe that the socle degrees should lie along polyhedrally describable subsets of \mathbb{Z}^d .

HISTORICAL CONTEXT

There have been a number of recent algorithmic computations in local cohomology, such as those by Walther (Wal99) (based on abstract methods of Lyubeznik (Lyu93)), Eisenbud–Mustață–Stillman (EMS00), Miller (Mil00), Mustață (Mus00), and Yanagawa (Yan02). These and related papers fall naturally into a number of categories. For instance,

the last three deal with \mathbb{Z}^d -graded modules over polynomial rings in d variables; in particular, they compute local cohomology with support on monomial ideals. In contrast, the paper (EMS00) works with coarser gradings—but still with monomial support, while (Wal99) requires no grading at all. As the gradings used become coarser, the papers increasingly depend on Gröbner bases: the monomial ideal papers require very little (if any) Gröbner basis computation; the coarser gradings depend heavily on commutative Gröbner bases; and the nongraded methods rely on noncommutative Gröbner bases over the Weyl algebra.

Regardless of the methods, all of the above papers share one fundamental aspect: the base ring is regular (usually a polynomial ring, in the algorithmic setting). The reason for restricting to these rings is that local cohomology over them behaves in many respects like a finitely generated module, even though it usually fails to be finitely generated. For example, Lyubeznik (Lyu93) and Walther (Wal99) take advantage of the fact that local cohomology modules over regular rings are finitely generated (indeed, holonomic) over the corresponding algebra of differential operators, and that the algebra of differential operators of a regular ring is easily presented, at least in characteristic zero.

Generally speaking, our methods lie somewhere between the monomial and coarsely graded methods described above, relying on a mix of Gröbner bases and integer programming. The principle underlying our computation of injective resolutions is that one should attempt to recover entire \mathbb{Z}^d -graded modules from their Q -graded parts. This idea originated for polynomial rings in (Mil98, Mus00, Mil00), was transferred in a restricted form to semigroup rings in (Yan01), and developed generally for semigroup-graded noetherian rings in (HM03). In the present context, the recovery of a module from its Q -graded part suggested that we compute injective resolutions via the irreducible resolutions of (Mil02).

Origins of the notion of sector partition can be seen in the Hilbert series formula for the local cohomology of canonical modules of normal semigroup rings (Ter99, Yan02), where the cellular homology was constant on large polyhedral regions of \mathbb{Z}^d . The accompanying notion of *straight module* (Yan01, HM03) abstracted this constancy; in fact, our Theorem 5.2 is really a theorem about straight modules as in (HM03, Definition 5.1). In any case, once the injective resolution has been computed using irreducible resolutions, the sector partition for local cohomology requires the entire \mathbb{Z}^d -graded structure of the injective resolution, and not just its Q -graded part.

CONVENTIONS AND NOTATION

In addition to the notation introduced thus far, we close this Introduction with a note on conventions. The semigroup Q is required to be saturated in Sections 5–7 because we do not know how to compute sector partitions in the unsaturated context (Remark 6.7). Other than the temporary saturation requirement in Section 3.2, the semigroup can be unsaturated in Sections 2–4. (Reminders of these conventions appear in each section).

The symbol $\mathbf{x}^\alpha \in \mathbb{k}[\mathbb{Z}^d]$ denotes a Laurent monomial in the localization $\mathbb{k}[\mathbb{Z}^d]$ of the semigroup ring $\mathbb{k}[Q]$. The \mathbb{k} -vector space spanned by $\{\mathbf{x}^\alpha \mid \alpha \in T\}$ for a subset $T \subseteq \mathbb{Z}^d$ will be denoted by $\mathbb{k}\{T\}$. The \mathbb{k} -subalgebra of $\mathbb{k}[\mathbb{Z}^d]$ will be denoted by $\mathbb{k}[T]$.

The **faces** of Q are those subsets minimizing linear functionals on Q . The **edges** and **facets** are the faces of dimension 1 and codimension 1. To every face F corresponds a prime ideal P_F and a quotient affine semigroup ring $\mathbb{k}[F] = \mathbb{k}\{F\}$.

All modules in this paper are \mathbb{Z}^d -graded unless otherwise stated. In particular, injective modules (defined in Section 5) are \mathbb{Z}^d -graded injective, which means that they are

usually not injective in the category of all $k[Q]$ -modules. Two subsets $S, T \subseteq \mathbb{Z}^d$ have the difference set $T - S = \{\beta - \alpha \mid \alpha \in S \text{ and } \beta \in T\} \subseteq \mathbb{Z}^d$. This allows us to write the localization of M along a face F as the module $M[\mathbb{Z}F] := M \otimes_{k[Q]} k[Q - F]$. Homomorphisms $N \rightarrow N'$ of modules are assumed to have \mathbb{Z}^d -graded degree $\mathbf{0}$, so that $N_\alpha \rightarrow N'_\alpha$ for all $\alpha \in \mathbb{Z}^d$.

We assume in this paper that standard algorithmic calculations with finitely generated modules over $k[Q]$ are available. In particular, we assume that the homology of any three-term (nonexact) sequence of finitely generated modules can be calculated, as can the submodule annihilated by a prime ideal of $k[Q]$. The \mathbb{Z}^d -grading only makes these computations easier, and the results of all such algorithms are still \mathbb{Z}^d -graded.

2. Effective irreducible hulls

In this section the affine semigroup Q need not be saturated. In the \mathbb{Z}^d -graded category \mathcal{M} from the Introduction, the injective modules have simple descriptions.

DEFINITION 2.1. Let $T \subset \mathbb{Z}^d$ be closed under addition of elements of $-Q$, by which we mean $T - Q \subset T$. Then $k\{T\}$ can be given the structure of a $k[Q]$ -module by setting

$$\mathbf{x}^a \mathbf{x}^\beta = \begin{cases} \mathbf{x}^{a+\beta} & \text{if } a + \beta \in T \\ 0 & \text{otherwise.} \end{cases}$$

An **indecomposable injective** is any module of the form $k\{\alpha + F - Q\}$, for some face F and $\alpha \in \mathbb{Z}^d$.

All such objects are injective in \mathcal{M} , and every injective object of \mathcal{M} is isomorphic to a finite direct sum of indecomposable injectives (MS03, Chapter 11). We shall work exclusively with objects in \mathcal{M} . Thus the term “injective module” in the rest of this paper will refer to modules of the above type.

Injectives are infinitely generated. For computations, we therefore work with certain finitely generated approximations. A module N is called **Q -graded** if N equals its **Q -graded part** $N_Q := \bigoplus_{a \in Q} N_a$. A submodule N of a module N' is an **essential submodule** if N intersects every nonzero submodule of N' nontrivially; the inclusion $N \hookrightarrow N'$ is also called an **essential extension**. In particular, N must be nonzero.

DEFINITION 2.2. An **irreducible sum** is a module that can be expressed as the Q -graded part J_Q of some injective module J . An **irreducible hull** of a Q -graded module N is an irreducible sum \overline{W} along with an essential extension $N \hookrightarrow \overline{W}$.

The existence of unique minimal injective resolutions (MS03, Corollary 11.35) includes the fact that every finitely generated module has an injective hull (that is, an inclusion into an injective that is an essential extension) that is unique up to isomorphism. Taking Q -graded parts yields immediately the following lemma.

LEMMA 2.3. *Every Q -graded module has an irreducible hull. It is unique up to isomorphism, and isomorphic to the Q -graded part of an injective hull of M .*

We call the modules of Definition 2.2 irreducible sums because of the next lemma, which is (Mil02, Lemma 2.2). An ideal W is called **irreducible** if W can not be expressed as an intersection of two ideals properly containing it.

LEMMA 2.4. *A monomial ideal W is irreducible if and only if the Q -graded part of some indecomposable injective module J satisfies $J_Q = \overline{W}$.*

Modules M are usually stored as data structures keeping track of their generators and relations—that is, as quotients of free modules. In the context of injective resolutions and local cohomology, storing M as a submodule of an irreducible sum is also useful. Our next definition specifies a data structure that precisely describes an irreducible sum \overline{W} .

DEFINITION 2.5. **Effective data** for an irreducible sum $\overline{W} = \bigoplus_{j=1}^r \mathbf{k}\{\alpha_j + F_j - Q\}_Q$ consist of:

1. an ordered r -tuple F_1, \dots, F_r of faces of Q ; and
2. an ordered r -tuple $\alpha_1, \dots, \alpha_r$, where $\alpha_j \in \mathbb{Z}^d / \mathbb{Z}F_j$ satisfies $Q \cap (\alpha_j + \mathbb{Z}F_j) \neq \emptyset$.

An **effective vector** of degree $a \in Q$ is an r -tuple $(\lambda_1, \dots, \lambda_r) \in \mathbf{k}^r$ such that $\lambda_j = 0$ whenever $a \notin \alpha_j + F_j - Q$. Concatenation of the respective face and degree data from two effective data yields their **direct sum**.

Note that the faces F_j need not be distinct. The condition $\alpha \in \mathbb{Z}^d / \mathbb{Z}F$ takes care of the fact that two degrees α and α' off by an element of $\mathbb{Z}F$ give the same module $\mathbf{k}\{\alpha + F - Q\} = \mathbf{k}\{\alpha' + F - Q\}$. Usually the α 's are recorded as elements of \mathbb{Z}^d , since the quotient mod $\mathbb{Z}F$ can be deduced from the face data. The condition $Q \cap (\alpha + \mathbb{Z}F) \neq \emptyset$ ensures that $\mathbf{k}\{\alpha + F - Q\}$ has nonzero Q -graded part. The condition on the λ 's simply requires each nonzero component to lie in a nonzero degree of the corresponding irreducible summand.

DEFINITION 2.6. An **effective irreducible hull** of a Q -graded module M consists of effective data for \overline{W} plus a list of finitely many effective vectors in \overline{W} generating a submodule isomorphic to M .

An irreducible hull $M \hookrightarrow \overline{W}$ is not quite dual to an expression $\mathcal{F} \rightarrow M$ as a quotient of a free module. The generators of \mathcal{F} have as their dual notion the face data F_1, \dots, F_r , which as abstract objects associated to M are known as **cogenerators**. Just as the degrees of the generators of \mathcal{F} need to be specified, so must the degree data for the cogenerators. However, the notion of effective vector for M as a submodule of \overline{W} is dual not to the notion of relation for M inside \mathcal{F} , but rather to the notion of cogenerator for M . Relations for M are, in actuality, dual to the notion of cogenerators for the *cokernel* of $M \hookrightarrow \overline{W}$, which correspond to indecomposable summands in cohomological degree 1 of the minimal injective resolution of M ; we dub these the **correlations** of M . Thus a presentation of M by generators and relations is dual to a presentation of M by cogenerators and correlations, whereas an irreducible hull presents M by generators and cogenerators.

3. Computing with irreducible hulls

Given a Q -graded module M in the usual way, via generators and relations, this section computes an irreducible hull $M \hookrightarrow \overline{W}$ as well as the cokernel of this inclusion.

Calculating an effective irreducible hull of M is, by definition, equivalent to calculating an irreducible decomposition of M . Thinking of the case $M = k[Q]/I$ for a monomial ideal I , this procedure is polyhedral in nature: it writes the set of monomials outside of I as a union of convex polyhedral regions whose facets are parallel to those of Q . The algorithm for computing an effective irreducible hull $M \hookrightarrow \overline{W}$, culminating in Proposition 3.7, does not require Q to be saturated.

Computing the cokernel, however, is strictly easier for saturated semigroups. The main point is the computation of generators for irreducible ideals. For saturated semigroups this is Proposition 3.14. The harder unsaturated case, in Proposition 3.16, relies on the computation of irreducible ideals over its saturation. To highlight the simplification in the saturated case, we state the main result of Sections 3.2–3.3 here.

PROPOSITION 3.1. *Generators and relations for M and \overline{W}/M are algorithmically computable from an effective irreducible hull $M \hookrightarrow \overline{W}$ over any affine semigroup ring $k[Q]$.*

PROOF. Generators for \overline{W} are already given, and relations for \overline{W} constitute a direct sum of irreducible ideals calculated as in Proposition 3.14 for saturated semigroups, and Proposition 3.16 in general. Since M is specified by its generators as a submodule of \overline{W} , the current proposition reduces to calculating submodules and quotients of modules presented by generators and relations. \square

3.1. EFFECTIVE IRREDUCIBLE HULLS FROM GENERATORS AND RELATIONS

This subsection does not require the affine semigroup Q to be saturated. The next two results make Algorithm 3.6 possible to state and easier to read. The notation $\langle y_1, \dots, y_j \rangle$ means ‘the $k[Q]$ -submodule generated by the elements y_1, \dots, y_j in their ambient module’, and $(0 :_M P_F)$ is the submodule of M annihilated by P_F .

LEMMA 3.2. *Suppose F has minimal dimension among faces of Q such that P_F is associated to M . Then the natural map $(0 :_M P_F)$ to its localization $(0 :_M P_F)[\mathbb{Z}F]$ along F is an inclusion. Furthermore, we can find algorithmically a set $B \subset (0 :_M P_F)$ of homogeneous elements that constitute a $k[\mathbb{Z}F]$ -basis for $(0 :_M P_F)[\mathbb{Z}F]$.*

PROOF. The $k[Q]$ -module $(0 :_M P_F)$ is naturally a torsion-free $k[F]$ -module, by minimality of $\dim F$. Therefore $(0 :_M P_F)$ includes into its localization along F , which must be a free $k[\mathbb{Z}F]$ -module. Now use the following algorithm. \square

ALGORITHM 3.3. (FOR LEMMA 3.2) Choose any element of $(0 :_M P_F)$ as the first basis vector $y_1 \in B$. Having chosen y_j , let y_{j+1} be any element of $(0 :_M P_F)$ whose image in $(0 :_M P_F)/\langle y_1, \dots, y_j \rangle$ generates a submodule of Krull dimension $\dim F$ (equivalently, the image of y_{j+1} has annihilator P_F). The algorithm terminates when the Krull dimension of the quotient $(0 :_M P_F)/\langle y_1, \dots, y_j \rangle$ is strictly less than $\dim F$.

LEMMA 3.4. *In the situation of Lemma 3.2, the scalar factor on the (monomial) coefficient of $y \in B$ in the unique $\mathbb{k}[\mathbb{Z}F]$ -linear combination of elements in B equaling any fixed element $z \in (0 :_M P_F)$ can be computed algorithmically.*

We present the proof as an algorithm.

ALGORITHM 3.5. (FOR LEMMA 3.4) Let $B(z) = \{y \in B \mid \deg(y) \equiv \deg(z) \pmod{\mathbb{Z}F}\}$. The coefficient of y in z is zero if $y \notin B(z)$. Otherwise, find elements a and $\{a_y \mid y \in B(z)\}$ in the face F such that $a + \deg(z) = a_y + \deg(y)$ for all $y \in B(z)$. By construction, $\{\mathbf{x}^{a_y} \cdot y \mid y \in B(z)\}$ is a \mathbb{k} -basis for the degree $a + \deg(z)$ piece of $(0 :_M P_F)$, and standard methods allow us to calculate the syzygy with $\mathbf{x}^a \cdot z$.

Write $\Gamma_F N := \Gamma_{P_F} N = (0 :_N P_F^\infty)$ for the set of elements in N annihilated by all high powers of P_F .

ALGORITHM 3.6.

```

INPUT    $Q$ -graded module  $M$  given by a generating set  $G \subset M$  and relations
OUTPUT  effective irreducible hull  $\overline{W}$  of  $M$  with effective vector set  $\Lambda$  indexed by  $G$ 
INITIALIZE  $N := M$ 
           $\overline{W} := (\{\}, \{\})$ , the empty effective datum for the irreducible sum 0
           $\lambda_g := ()$  for all  $g \in G$ ; here  $()$  is the effective vector of length zero in  $\overline{W}$ 
           $i := 1$ 
DEFINE   $(F_1, \dots, F_s) :=$  an ordering of the faces of  $Q$  with  $\dim(F_i) \leq \dim(F_{i+1})$ 
WHILE    $i \leq s$  DO
  DEFINE  $F := F_i$ 
           $B := \mathbb{k}[\mathbb{Z}F]$ -basis for  $(0 :_N P_F)[\mathbb{Z}F]$ , as in Algorithm 3.3
  WHILE  $y \in B$  and  $g \in G$  DO
    IF    $(0 :_{\langle g \rangle} P_F) \neq 0$  in some degree  $a_{yg} \equiv \deg(y) \pmod{\mathbb{Z}F}$ 
      THEN  $\lambda_{yg} :=$  scalar coefficient of  $y$  on  $\mathbf{x}^{a_{yg} - \deg(g)} \cdot g$ , as in Algorithm 3.5
      ELSE  $\lambda_{yg} := 0$ 
    END IF-THEN-ELSE
  END WHILE-DO
  REDEFINE  $\lambda_g :=$  concatenation of the two vectors  $\lambda_g$  and  $(\lambda_{yg})_{y \in B}$ , for  $g \in G$ 
           $\overline{W} := \overline{W} \oplus (\#B \text{ copies of } F, \mathbb{Z}^d\text{-degrees of vectors in } B)$ 
           $N := M/\Gamma_F M$ 
           $i := i + 1$ 
END WHILE-DO
OUTPUT   $\overline{W}$  along with  $\Lambda = \{\lambda_g\}_{g \in G}$ , where  $\lambda_g$  is in degree  $\deg(g)$ 

```

PROPOSITION 3.7. *Algorithm 3.6 outputs an effective irreducible hull of M , using generators and relations for M as input.*

PROOF. We must show that the homomorphism $M \rightarrow \overline{W}$ determined by G and Λ is well-defined and injective. More precisely: monomial combinations z of the generators of

M are zero if and only if the corresponding monomial combinations z_λ of the λ_g are zero in \overline{W} ; here, λ_g represents not a data structure but an element of \overline{W} .

The combination z is nonzero in M if and only if the submodule $\langle z \rangle \subseteq M$ generated by z has an associated prime. The associated prime is $F := F_i$ if and only if the image of $\langle z \rangle$ in the successive quotient $N = M/\Gamma_{F_{i-1}}M$ intersects $(0 :_N P_F)$ nontrivially (this in particular implies that $(0 :_N P_F)$ is nonzero, so P_F is associated to M). This nontriviality of $\langle z \rangle \cap (0 :_N P_F)$ is equivalent to having at least one of the terms monomial $\cdot g$ appearing in z be nonzero in the same $(0 :_N P_F)$, because B is a basis for $(0 :_N P_F)[\mathbb{Z}F]$. Finally, monomial $\cdot g$ is nonzero precisely when the corresponding element monomial $\cdot \lambda_g$ has nonzero coefficient in the appropriate summand of \overline{W} . \square

REMARK 3.8. Some alterations to Algorithm 3.6 may improve its running time.

1. It is possible to avoid taking the successive quotients $N/\Gamma_F M$ at the REDEFINE step. These quotients are designed to make Lemmas 3.2 and 3.4 apply, as well as to make N successively simpler. However, the cost of taking these quotients may not be worth it, since the final sentence of Lemma 3.2 holds even if F doesn't have minimal dimension (so $(0 :_M P_F)$ doesn't include into its localization along F). In fact, both of Algorithms 3.3 and 3.5 still work in this more general setting.
2. Of the faces on the list (F_1, \dots, F_s) , only those associated to M need to be tested. If desired, these faces can be detected using homological methods.
3. Instead of computing and working with $(0 :_M P_F)$ for each face separately, one could work with the modules $(0 :_M I_c)$ for each c , where I_c is the intersection of all primes P_F for faces F of dimension c .

3.2. GENERATORS AND RELATIONS FROM IRREDUCIBLE HULLS: SATURATED CASE

In this subsection we assume that Q is saturated. Our goal is to compute relations on the generators for M that come as part of an effective irreducible hull $M \hookrightarrow \overline{W}$. As we shall see in the proof of Proposition 3.1, the computation essentially reduces to the case where $M = \overline{W}$ is an *indecomposable* irreducible sum \overline{W} , so we are to determine the kernel of the surjection $\mathbf{k}[Q] \rightarrow \overline{W}$. More explicitly, given a face F and a degree $a \in Q$, we must find generators of

$$W := \mathbf{k}\{Q \setminus (a + F - Q)\} \quad (3.1)$$

as an ideal in $\mathbf{k}[Q]$.

Since Q is saturated, there is a unique minimal set of oriented hyperplanes inside \mathbb{Z}^d whose closed positive half-spaces in \mathbb{Z}^d have intersection equal to Q . The map sending $H \mapsto H \cap Q$ gives a bijection from these hyperplanes to the facets of Q . Denote by H_+ the closed positive half-space determined by an oriented hyperplane H , and by H_+° the *open* positive half-space. Thus H_+° is the complement of $-H_+$ but can also be characterized as the lattice distance 1 translate of H_+ in the positive direction.

LEMMA 3.9. *Given any face F of Q and any element $a \in Q$,*

$$Q \setminus (a + F - Q) = \bigcup_{H \supseteq F} (a + H_+^\circ) \cap Q.$$

PROOF. We have $a + F - Q = \bigcap_{H \supseteq F} a - H_+$ because Q is saturated (recall $F - Q = -(Q + \mathbb{Z}F)$). Thus $\mathbb{Z}^d \setminus (a + F - Q) = \bigcup_{H \supseteq F} a + H_+^\circ$. Now intersect with Q . \square

Lemma 3.9 reduces the computation of generators for W as in (3.1) to the case where F is itself a facet, at least when Q is saturated. The next algorithm and two lemmas cover this case by producing some rational polytopes whose integer points do the job. For notation, \mathbb{R}_+F denotes the real cone generated by F in $\mathbb{R}^d = \mathbb{R} \otimes \mathbb{Z}^d$, and $\mathbb{R}H$ denotes the real span of a hyperplane H . Also, by a Q -set we mean a subset of \mathbb{Z}^d closed under addition by elements of Q . A set G of vectors in \mathbb{Z}^d **generates** a Q -set T if $T = G + Q$.

LEMMA 3.10. *Let G_Q be the zonotope that is the Minkowski sum of all primitive integer vectors along rays of Q . Then, for all $\alpha \in \mathbb{R}^d$, the lattice points in $\alpha + G_Q$ generate $(\alpha + \mathbb{R}_+Q) \cap \mathbb{Z}^d$ as a Q -set.*

PROOF. Let β be a lattice point in $\alpha + \mathbb{R}_+Q$. If there is no primitive integer vector ρ along a ray of Q such that $\beta - \rho$ still lies in $\alpha + \mathbb{R}_+Q$, then $\beta \in \alpha + G_Q$. \square

ALGORITHM 3.11.

```

INPUT    $Q :=$  a saturated semigroup
         $H :=$  one of the hyperplanes bounding  $Q$ 
         $a \in Q$ 
OUTPUT  finite set  $B \subset Q$  such that the ideal  $\langle \mathbf{x}^b \mid b \in B \rangle$  equals  $\mathfrak{k}\{(a + H_+^\circ) \cap Q\}$ 
DEFINE   $G :=$  the polytope  $G_Q$  in Lemma 3.10
         $F := H \cap Q$ , a facet of  $Q$ 
         $\Delta :=$  the set of faces of  $Q$  intersecting  $F$  only at  $\mathbf{0} \in Q$ 
INITIALIZE  $B := \{\}$ , the empty subset of  $Q$ 
WHILE    $D \in \Delta$  DO
    DEFINE  $B_D :=$  lattice points in Minkowski sum  $((a + \mathbb{R}H) \cap \mathbb{R}_+D) + G$ 
    REDEFINE  $B := B \cup B_D$ 
    NEXT  $D$ 
END     WHILE-DO
OUTPUT  $B$ 

```

LEMMA 3.12. *Algorithm 3.11 computes generators for the ideal $\mathfrak{k}\{(a + H_+^\circ) \cap Q\}$.*

PROOF. Suppose $b \in (a + H_+^\circ) \cap Q$. The intersection $(b + \mathbb{R}H) \cap \mathbb{R}_+Q$ is a polyhedron whose bounded faces are precisely the polytopes $(b + \mathbb{R}H) \cap \mathbb{R}_+D$ for $D \in \Delta$, and whose recession cone is \mathbb{R}_+F . Therefore $b \in b' + \mathbb{R}_+F$ for some real vector $b' \in (b + \mathbb{R}H) \cap \mathbb{R}_+D$ and some face $D \in \Delta$. Moreover, b' lies in $b'' + \mathbb{R}_+D$ for some real vector $b'' \in (a + \mathbb{R}H) \cap \mathbb{R}_+D$. Consequently, b lies in $b'' + \mathbb{R}_+(D + F)$, and therefore in $b'' + \mathbb{R}_+Q$. Now \mathbf{x}^b lies in the $\mathfrak{k}[Q]$ -module generated by $\mathfrak{k}\{B_D\}$, by definition of G . \square

REMARK 3.13. Some alterations to Algorithm 3.11 may improve its running time.

1. Instead of computing just one polytope $G = G_Q$ and Minkowski summing it to define every B_D , we could define B_D with G_{D+F} in place of G , for each face $D \in \Delta$.

This might reduce the number of lattice points in B dramatically, but would require more computations as in Lemma 3.10.

2. Restricting to the maximal elements in Δ will speed things up.

Let us summarize the above algorithm and three lemmas. (See Section 7 for issues concerning the output of the algorithm in the following proposition, and post-processing for the purpose of reducing its complexity.)

PROPOSITION 3.14. *Generators of the irreducible ideal $W = \ker(\mathbb{k}[Q] \rightarrow \overline{W})$ are algorithmically computable using as input an indecomposable effective irreducible sum \overline{W} over a normal semigroup ring $\mathbb{k}[Q]$.*

PROOF. Apply Algorithm 3.11 to each of the sets $(a + H_+^\circ) \cap Q$ in Lemma 3.9. \square

3.3. GENERATORS AND RELATIONS FROM IRREDUCIBLE HULLS: UNSATURATED CASE

Now we return to the general case, where Q need not be saturated, and denote by Q^{sat} the saturation of Q . The basic idea for computing generators of irreducible ideals in $\mathbb{k}[Q]$ is to intersect (as $\mathbb{k}[Q]$ -modules) the submodule $\mathbb{k}[Q] \subset \mathbb{k}[Q^{\text{sat}}]$ with the ideal $W \subseteq \mathbb{k}[Q^{\text{sat}}]$ output in the saturated case, Proposition 3.14. Then it remains to find the appropriate F -primary component of W as a $\mathbb{k}[Q]$ -module, where F is the unique face of dimension $\dim(\overline{W})$ associated to \overline{W} (as a $\mathbb{k}[Q]$ -module).

Every module in Algorithm 3.15 is to be considered as a $\mathbb{k}[Q]$ -module—even those generated as $\mathbb{k}[Q^{\text{sat}}]$ -modules. Thus F is always a face of Q , and we consider $F - Q$ as opposed to $F - Q^{\text{sat}}$. Note, however, that $\mathbb{k}\{F - Q^{\text{sat}}\}$ does equal the corresponding injective over $\mathbb{k}[Q^{\text{sat}}]$, even though F is a face of Q ; subtracting Q^{sat} automatically saturates F .

ALGORITHM 3.15.

```

INPUT    $Q :=$  a semigroup, not necessarily saturated
         $F :=$  a face of  $Q$ 
         $a \in Q$ 
OUTPUT   $B \subset Q$  such that  $\langle \mathbf{x}^b \mid b \in B \rangle$  equals the ideal  $\mathbb{k}\{Q \setminus (a + F - Q)\}$  in  $\mathbb{k}[Q]$ 
DEFINE   $\overline{V} := \mathbb{k}\{Q^{\text{sat}} \setminus (a + F - Q^{\text{sat}})\}$ , an indecomposable irreducible over  $\mathbb{k}[Q^{\text{sat}}]$ 
         $V :=$  the kernel of  $\mathbb{k}[Q^{\text{sat}}] \rightarrow \overline{V}$  output by Proposition 3.14
         $W := V \cap \mathbb{k}[Q]$ , the intersection taken inside  $\mathbb{k}[Q^{\text{sat}}]$ 
         $I := \bigcap \{P_D \mid D \text{ is a facet of } F\}$ , an ideal in  $\mathbb{k}[Q]$ 
INITIALIZE  $B :=$  degrees of the elements generating  $W$ 
         $\overline{W} := \mathbb{k}[Q]/W$ 
WHILE    $(0 :_{\overline{W}} P_F)$  has a generator in some degree  $\not\equiv a \pmod{\mathbb{Z}F}$  DO
    DEFINE  $G :=$  generators for  $(0 :_{\overline{W}} P_F)$  that lie in degrees  $\not\equiv a \pmod{\mathbb{Z}F}$ 
    REDEFINE  $B := B \cup$  degrees of the elements in  $G$ 
             $\overline{W} := \overline{W}/G$ 
    DEFINE  $G' :=$  generators for  $\Gamma_I \overline{W}$ 
    REDEFINE  $B := B \cup$  degrees of the elements in  $G'$ 
             $\overline{W} := \overline{W}/G'$ 
END     WHILE-DO

```

OUTPUT B

PROPOSITION 3.16. *Algorithm 3.15 outputs generating degrees for $\mathbf{k}\{Q \setminus (a + F - Q)\}$.*

PROOF. The module \overline{W} gets initialized as a quotient of $\mathbf{k}[Q^{\text{sat}}]$ with dimension $\dim(F)$ as a $\mathbf{k}[Q]$ -module. This much holds by the saturated version Proposition 3.14 applied to \overline{V} , and the preservation of dimension (Eis95, Proposition 9.2) for the module-finite ring extension $\mathbf{k}[Q] \subseteq \mathbf{k}[Q^{\text{sat}}]$ applied to V . One part of the output is clear: the set $\langle \mathbf{x}^b \mid b \in B \rangle$ generates the kernel of the map $\mathbf{k}[Q] \rightarrow \overline{W}$ at every stage in the algorithm. The question is whether \overline{W} is the claimed indecomposable irreducible sum.

In the first REDEFINE step, the annihilator of $\mathbf{x}^a \in \overline{W}$ remains P_F . Indeed, any element killed by P_F that generates a submodule containing a nonzero element in degree a must itself have degree congruent to $a \pmod{\mathbb{Z}F}$. The second REDEFINE step only kills elements with annihilators strictly larger than that of \mathbf{x}^a ; such elements can't generate submodules containing \mathbf{x}^a . Therefore, \overline{W} has only one associated prime P_F after each loop of WHILE-DO, by dimension considerations.

When the loop terminates, the localization $(0 :_{\overline{W}} P_F)[\mathbb{Z}F]$ along F is indecomposable, being isomorphic to $\mathbf{k}\{a + \mathbb{Z}F\}$. It follows that the kernel of the surjection $\mathbf{k}[Q] \rightarrow \overline{W}$ is an irreducible ideal (Vas98, Proposition 3.1.7). We are done by Lemma 2.4, because $\mathbf{k}\{a + F - Q\}_Q$ is the only indecomposable irreducible sum for which the annihilator of \mathbf{x}^a is P_F . \square

REMARK 3.17. Some alterations to Algorithm 3.15 may improve its efficiency.

1. The step $\overline{W} := \overline{W}/\Gamma_I \overline{W}$ need not occur until the very last step before OUTPUT. Its current placement is designed to speed the computation by simplifying \overline{W} in each loop, but the cost of taking the colon may not make up for it. Instead, the end of the algorithm can be replaced by:

```

WHILE (0 :_{\overline{W}} P_F) has rank strictly larger than 1 over \mathbf{k}[F] DO
  DEFINE G := generators for (0 :_{\overline{W}} P_F) lying in degrees \not\equiv a \pmod{\mathbb{Z}F}
  REDEFINE B := B \cup degrees of the elements in G
           \overline{W} := \overline{W}/G
END WHILE-DO
REDEFINE B := B \cup degrees of the generators of \Gamma_I \overline{W}
OUTPUT B

```

2. As in Remark 3.8, it is not necessary to compute all of $(0 :_{\overline{W}} P_F)$ in the WHILE-DO loop. It suffices instead to let G be a basis for $(0 :_{\overline{W}} P_F)[\mathbb{Z}F]$. This remark also holds for the reworked WHILE-DO loop in the previous item.
3. The set B can become rather redundant. Since the machine will have to keep a presentation of \overline{W} in memory, the algorithm could simply spit out the relations defining \overline{W} as a $\mathbf{k}[Q]$ -module at the very end, without keeping track of B at all.

4. Computing injective resolutions

In this section the semigroup Q is not required to be saturated. Our goal is the main result (Theorem 4.7) in the first half of the paper: an algorithm to compute injective

resolutions of finitely generated modules over $k[Q]$, in the \mathbb{Z}^d -graded setting. That is, given generators and relations for a finitely generated \mathbb{Z}^d -graded module M , we will compute an exact sequence $0 \rightarrow M \rightarrow J^0 \rightarrow J^1 \rightarrow \dots$ in which J^i is a \mathbb{Z}^d -graded injective module for each i . Of course, we shall only say how to calculate up to some specified cohomological degree, as injective resolutions usually do not terminate. This will not pose a problem for our subsequent computation in Section 7 of local cohomology, which vanishes past cohomological degree $d + 1$ anyway.

The upshot is to reduce the computation of injective resolutions to finding irreducible hulls of finitely generated Q -graded modules and computing their cokernels, which we have already done in Section 3.

The data structures we employ for \mathbb{Z}^d -graded injective resolutions are the matrices we introduce in the next definition.

DEFINITION 4.1. A **monomial matrix** is a matrix of constants λ_{qp} along with

1. a vector $\alpha_q \in \mathbb{Z}^d$ and a face $F_q \in Q$ for each row, and
2. a vector $\alpha_p \in \mathbb{Z}^d$ and a face $F_p \in Q$ for each column

such that $\lambda_{qp} = 0$ unless $F_p \subseteq F_q$ and $\alpha_p \in \alpha_q + F_q - Q$.

These monomial matrices generalize those in (Mil00), which were for $Q = \mathbb{N}^d$.

To any monomial matrix we can associate a map $J \mapsto J'$ of injective modules in the following manner. Each row and column label gives the data of an indecomposable injective; we think of the row labels as giving summands of J and the column labels as giving summands of J' . To give a map from J to J' is thus the same as giving a matrix of maps from the row indecomposables to the column indecomposables. Such a map $k\{\alpha_q + F_q - Q\} \mapsto k\{\alpha_p + F_p - Q\}$ is necessarily zero unless $F_p \subseteq F_q$ and $\alpha_p \in \alpha_q + F_q - Q$. In the latter case it is determined by a single scalar λ_{qp} . Hence

$$\begin{array}{c}
 \dots F_p \dots \\
 \dots \alpha_p \dots \\
 \vdots \quad \vdots \\
 F_q \alpha_q \left(\begin{array}{c} \lambda_{qp} \end{array} \right) \\
 \vdots \quad \vdots
 \end{array}$$

is a monomial matrix representing a map

$$\bigoplus_q k\{\alpha_q + F_q - Q\} \mapsto \bigoplus_p k\{\alpha_p + F_p - Q\}.$$

The component $k\{\alpha_q + F_q - Q\} \mapsto k\{\alpha_p + F_p - Q\}$ of this homomorphism takes \mathbf{x}^α to $\lambda_{qp}\mathbf{x}^\alpha$ for all $\alpha \in \alpha_p + F_p - Q$, and is zero elsewhere.

Note that in degree α , the map $J_\alpha \mapsto J'_\alpha$ given by a monomial matrix is obtained by deleting the rows and columns labeled by α_p, F_p such that α does not lie in $\alpha_p + F_p - Q$. (This corresponds to ignoring those summands of J and J' not supported at α .) Ignoring the labels on what remains gives us a matrix with entries in k , which defines the k -vector space map $J_\alpha \mapsto J'_\alpha$.

Two monomial matrices represent the same map of injectives (with given decompositions into direct sums of indecomposable injectives) if and only if (i) their scalar entries

are equal, (ii) the corresponding faces F_r are equal, where $r = p, q$, and (iii) the corresponding vectors α_r are congruent modulo $\mathbb{Z}F_r$.

Rather than compute directly with cumbersome, infinitely generated injectives, it is more convenient to approximate injective resolutions using irreducible sums.

DEFINITION 4.2. An **irreducible resolution** of a Q -graded module M is an exact sequence $0 \rightarrow M \rightarrow \overline{W}^0 \rightarrow \overline{W}^1 \rightarrow \dots$ in which each \overline{W}^j is an irreducible sum.

Irreducible resolutions are approximations to injective resolutions; indeed, the Q -graded part of any injective resolution is an irreducible resolution (Mil02, Theorem 2.4). In particular, monomial matrices just as well represent homomorphisms of irreducible sums, as long as the degree labels α_q and α_p all can be chosen to lie in Q . The (apparent) advantage to irreducible resolutions over injective resolutions is their finiteness.

COROLLARY 4.3. *For any finitely generated Q -graded $k[Q]$ -module M , Propositions 3.1 and 3.7 inductively compute a minimal irreducible resolution \overline{W}^\bullet of M algorithmically.*

PROOF. Minimal irreducible resolutions have finite length (that is, they vanish in all sufficiently high cohomological degrees) by (Mil02, Theorem 2.4). The computability therefore follows from Propositions 3.1 and 3.7 by induction on the highest cohomological degree required. \square

The next result demonstrates the precise manner in which irreducible resolutions approximate injective resolutions for computational purposes.

PROPOSITION 4.4. *Let M be a finitely generated module with minimal injective resolution J^\bullet and minimal irreducible resolution \overline{W}^\bullet . Suppose that every indecomposable summand in the first n cohomological degrees of J^\bullet has nonzero Q -graded part. Then M is Q -graded, and the data contained in the first n stages of \overline{W}^\bullet constitute a finite data structure for the first n cohomological degrees of J^\bullet .*

PROOF. Every map in J^\bullet can be expressed using the finite data of a monomial matrix, and this data can be read immediately off the maps in \overline{W}^\bullet . \square

If we can algorithmically determine a \mathbb{Z}^d -graded shift of M so that the hypotheses of Proposition 4.4 are satisfied, then we can compute the minimal injective resolution of M up to cohomological degree n . This task requires a lemma, in which \mathfrak{m} denotes the maximal ideal $P_{\{0\}}$ generated by nonunit monomials in $k[Q]$.

LEMMA 4.5. *Let J^\bullet be a minimal injective resolution of a finitely generated module M , and F a face of Q . If every indecomposable summand of $\Gamma_{\mathfrak{m}}J^{j+d-\dim(F)}$ has nonzero Q -graded part, then every indecomposable summand of J^j isomorphic to a \mathbb{Z}^d -graded shift of $k\{F - Q\}$ has nonzero Q -graded part.*

PROOF. (HM03, Proposition 3.5), in the special case of an affine semigroup ring. \square

Every indecomposable summand of $\Gamma_{\mathfrak{m}}J^j$ is a shift $k\{\alpha - Q\}$ of $k\{-Q\}$. Such an indecomposable injective has nonzero Q -graded part if and only if $\alpha \in Q$. Our final

lemma in this section describes the (standard) way to calculate the shifts α appearing in $\Gamma_{\mathfrak{m}} J^j$. The number $\mu^{j,\alpha}(M)$ of shifts $\mathfrak{k}\{\alpha - Q\}$ appearing as summands in cohomological degree j of the minimal injective resolution of M is called the j^{th} **Bass number** of M in degree α .

LEMMA 4.6. *Let \mathcal{F}^\bullet be a free resolution of the residue field \mathfrak{k} . The Bass number $\mu^{j,\alpha}(M)$ is effectively computable as the \mathfrak{k} -vector space dimension of $H^j(\mathcal{F}^\bullet, M)_\alpha$.*

PROOF. This expression of Bass numbers as dimensions (over \mathfrak{k}) of Ext modules is standard; see (BH93, Chapter 3). The computability follows because we can calculate free resolutions, homomorphisms, and homology over $\mathfrak{k}[Q]$. \square

Now we come to our central result. For notation, $M(-a)$ denotes the \mathbb{Z}^d -graded shift of M up by a , so that $M(-a)_b = M_{b-a}$.

THEOREM 4.7. *Fix a finitely generated $\mathfrak{k}[Q]$ -module M and an integer i . There is an algorithmically computable $a \in Q$ for which Propositions 3.1 and 3.7 inductively compute the minimal injective resolution of $M(-a)$ through cohomological degree $i + 1$.*

PROOF. After using Lemma 4.6 to compute the Bass numbers of M up to cohomological degree $i + 1 + \dim(M)$, choose a so that the corresponding Bass numbers of $M(-a)$ have \mathbb{Z}^d -graded degrees lying in Q . At this point, $M(-a)$ satisfies the hypotheses of Proposition 4.4 with $n = i + 1$, by Lemma 4.5. Now apply Corollary 4.3. \square

5. Sector partitions from injectives

We turn now to sector partitions, for which we assume henceforth that the affine semi-group Q is saturated. As a prerequisite to producing sector partitions of local cohomology modules, we demonstrate in this section that injective modules admit sector partitions, as does the homology of any complex of injective modules.

PROPOSITION 5.1. *Suppose $J = \bigoplus_{i=1}^r J_i$ is an injective module decomposed into summands $J_i = \mathfrak{k}\{\alpha_i + F_i - Q\}$. For each subset $A \subseteq \{1, \dots, r\}$ define S_A to be the set*

$$S_A = \{ \alpha \in \mathbb{Z}^d \mid (J_i)_\alpha \cong \mathfrak{k} \text{ for } i \in A \}$$

of all degrees in \mathbb{Z}^d such that the summands of J nonzero in that degree are precisely those indexed by A . The sets S_A canonically determine a sector partition $\mathcal{S}(J) \vdash J$.

PROOF. For each $\alpha \in \mathbb{Z}^d$, either $(J_i)_\alpha = \{0\}$ or $(J_i)_\alpha = \mathfrak{k} \cdot \mathbf{x}^\alpha$. Therefore $\mathcal{S}(J)$ is indeed a partition of \mathbb{Z}^d . Now we must show that S_A is a finite union of polyhedra as in part 1 of Definition 1.2. The set $\alpha + F - Q$ of degrees is the set of lattice points in a polyhedron of the desired form because the half-spaces whose intersection is $\alpha + F - Q$ are bounded by hyperplanes parallel to facets of Q , by definition. These hyperplanes divide \mathbb{Z}^d into finitely many disjoint regions (place the lattice points lying on each hyperplane in the region on the positive side of that hyperplane), each of which consists of the lattice points in a polyhedron of the desired form. Thus the complement $\mathbb{Z}^d \setminus (\alpha + F - Q)$ is the required kind of finite union. We conclude that S_A is a finite union of regions, each of which is an intersection of r polyhedral regions—one from each of the summands J_i .

For each index set A such that S_A is nonempty, define $J_{S_A} \subseteq k^r$ to be the subspace spanned by the basis vectors e_i such that $i \in A$. Then for each degree α in S_A , the map $J_\alpha \rightarrow J_{S_A}$ required by part 2 of Definition 1.2 can be taken to equal the zero map on $(J_i)_\alpha$ for i not in A , and the map sending \mathbf{x}^α to e_i on $(J_i)_\alpha$ for i in A .

To define the maps $\mathbf{x}^{S_B - S_A}$ for index sets A and B such that $S_B - S_A$ is nonempty, as in part 3 of Definition 1.2, it suffices to define the image of e_i for each i in A . We take $\mathbf{x}^{S_B - S_A}(e_i) = e_i$ if i is in B , and $\mathbf{x}^{S_B - S_A}(e_i) = 0$ otherwise. Commutativity of the required diagram follows from the definition of the module structure on $k\{\alpha_i + F_i - Q\}$. Specifically, for $\alpha \in S_A$ and $\beta \in S_B$ with $\beta - \alpha \in Q$, multiplication by $\mathbf{x}^{\beta - \alpha}$ takes \mathbf{x}^α to \mathbf{x}^β in J_i for $i \in B$, and takes \mathbf{x}^α to zero in J_i for i outside B . \square

The sector partition in Proposition 5.1 descends to the cohomology H of any complex of injectives, via monomial matrices. The forthcoming sector partition of H is really determined *canonically* by J^\bullet (without its direct sum decomposition), even though the way we present things here makes it look like bases must be chosen. We chose this route because bases are good for computation, while uniqueness is immaterial.

THEOREM 5.2. *If H is a module that can be expressed as the (middle) homology of a complex $J^\bullet : J' \rightarrow J \rightarrow J''$ in which all three modules are injective, or all three modules are flat, then there is a sector partition $\mathcal{S}(J^\bullet) \vdash H$ determined by J^\bullet .*

PROOF. Choose direct sum decompositions to write

$$J' = \bigoplus_{i=1}^{r'} J'_i, \quad J = \bigoplus_{i=1}^r J_i, \quad \text{and} \quad J'' = \bigoplus_{i=1}^{r''} J''_i.$$

Let Φ and Ψ be the monomial matrices representing the maps $J' \mapsto J$ and $J \mapsto J''$, respectively. The sectors in the sector partition $\mathcal{S}(J' \oplus J \oplus J'') \vdash J' \oplus J \oplus J''$ are indexed by triples (A', A, A'') of subsets of $\{1, \dots, r'\}$, $\{1, \dots, r\}$, $\{1, \dots, r''\}$, respectively, and automatically satisfy the polyhedrality condition in part 1 of Definition 1.2 by Proposition 5.1. We take $\mathcal{S}(J^\bullet)$ to partition \mathbb{Z}^d into these sectors.

For each triple (A', A, A'') we have maps $\Phi_{A'}^A : J_{S_{A'}}^A \rightarrow J_{S_A}$ and $\Psi_A^{A''} : J_{S_A} \rightarrow J_{S_{A''}}^{A''}$, whose monomial matrices are defined by deleting: row i' of Φ for i' not in A' ; column i of Φ and row i of Ψ for i not in A ; and column i'' of Ψ for i'' not in A'' . Let

$$H_{S_{A'}, A, A''} = \ker(\Psi_A^{A''}) / \text{im}(\Phi_{A'}^A). \quad (5.1)$$

For any α in $S_{A', A, A''}$, we have a commutative diagram

$$\begin{array}{ccccc} J'_\alpha & \longrightarrow & J_\alpha & \longrightarrow & J''_\alpha \\ \downarrow & & \downarrow & & \downarrow \\ J'_{S_{A'}} & \xrightarrow{\Phi_{A'}^A} & J_{S_A} & \xrightarrow{\Psi_A^{A''}} & J''_{S_{A''}} \end{array} \quad (5.2)$$

that induces the required isomorphism $H_\alpha \cong H_{S_{A'}, A, A''}$. It is routine to check that the maps $H_{S_{A'}, A, A''} \rightarrow H_{S_{B'}, B, B''}$ induced from the corresponding maps on $J'_{A'}$, J_A , and $J''_{A''}$ commute with this isomorphism. \square

Once we have Theorem 5.2, the only step remaining to prove Theorem 1.3 is to exhibit $H_i^i(M)$ as the homology of a complex of injectives.

REMARK 5.3. The results in this section hold just as well for **flat** objects of \mathcal{M} , which are Matlis dual to injective objects and hence isomorphic to finite direct sums of modules of the form $k\{\alpha + F + Q\}$ for some α in \mathbb{Z}^d and some face F of Q (MS03, Chapter 11). For the proofs, simply apply Matlis duality to the results for injectives.

6. Computing sector partitions

Again letting Q be a saturated affine semigroup, the next task is actually computing the finitely many polyhedra whose lattice points comprise the sectors in the sector partition $\mathcal{S}(J) \vdash J$ of an injective module. That is, we need to make Proposition 5.1 and its proof into an algorithm.

Since Q is saturated, there are unique primitive integer linear functionals τ_1, \dots, τ_n taking $\mathbb{Z}^d \rightarrow \mathbb{Z}$, one for each facet of Q , such that $Q = \bigcap_{i=1}^n \{\tau_i \geq 0\}$ is the set of lattice points in the intersection of their positive half-spaces. The degrees on which indecomposable injectives are supported can be expressed in terms of these linear functionals, via the following identity:

$$\alpha + F - Q = \{\beta \in \mathbb{Z}^d \mid \tau_i(\beta) \leq \tau_i(\alpha) \text{ whenever } \tau_i(F) = 0\}. \quad (6.1)$$

In other words, $F - Q$ is the intersection of the *negative* half-spaces for those functionals τ_i vanishing on F , and $\alpha + F - Q$ is simply a translate. By convention, we use the notation $\tau_i(\beta) \leq \infty$ to mean that there is no restriction on the value of $\tau_i(\beta)$. This allows a notation $\tau_F(\alpha) \in (\mathbb{Z} \cup \infty)^n$ for the vector whose i^{th} coordinate satisfies

$$\tau_F(\alpha)_i = \begin{cases} \tau_i(\alpha) & \text{if } \tau_i(F) = 0 \\ \infty & \text{otherwise.} \end{cases}$$

The point is that a vector $\beta \in \mathbb{Z}^d$ lies in $\alpha + F - Q$ if and only if $\tau(\beta) \leq \tau_F(\alpha)$, where

$$\tau(\beta) = (\tau_1(\beta), \dots, \tau_n(\beta))$$

and the ‘ \leq ’ symbol denotes componentwise comparison. We shall use the corresponding definitions of $\tau_F(\alpha)$ and $\tau(\beta)$ for vectors $\alpha, \beta \in \mathbb{R}^d = \mathbb{R} \otimes \mathbb{Z}^d$, so $\tau_F(\alpha) \in (\mathbb{R} \cup \infty)^n$.

For the rest of this section, let

$$J = \bigoplus_{j=1}^r J^j, \quad \text{with } J^j = k\{\alpha_j + F_j - Q\}, \quad (6.2)$$

be an injective module, and define

$$\tau^j := \tau_{F_j}(\alpha_j) \quad \text{for } j = 1, \dots, r.$$

Thus for $i = 1, \dots, n$ the vector τ^j has i^{th} coordinate $\tau_i^j = \tau_{F_j}(\alpha_j)_i$, which equals either $\tau_i(\alpha_j)$ or ∞ , depending on whether τ_i vanishes on F_j or not. Even without calculating the set $\mathcal{S}(J)$ algorithmically, the vectors τ^j specify the map from \mathbb{Z}^d to $\mathcal{S}(J)$, by definition. We record a precise version of this statement in the next lemma.

LEMMA 6.1. *A degree $\alpha \in \mathbb{Z}^d$ lies in S_A if and only if $A = \{j \in \{1, \dots, r\} \mid \tau(\alpha) \leq \tau^j\}$.*

It remains to ascertain which sets S_A of lattice points are nonempty, and to determine the pairs A, B for which we must compute a map $\mathbf{x}^{B-A} : J_A \rightarrow J_B$. (The maps themselves, which are canonical, are constructed in the proof of Proposition 5.1.) For each

functional τ_i there is a permutation w_i of $\{1, \dots, r\}$ satisfying $\tau_i^{w_i(1)} \leq \dots \leq \tau_i^{w_i(r)}$. To simplify notation, we write $\tilde{\tau}_i^\ell$ instead of $\tau_i^{w_i(\ell)}$. Also, set $\tilde{\tau}_i^0 = -\infty$ and $\tilde{\tau}_i^{r+1} = \infty$.

For fixed i , the parallel affine hyperplanes $\{\tau_i = \tilde{\tau}_i^\ell\}_{\ell=1}^r$ divide \mathbb{Z}^d into strips

$$\{\beta \in \mathbb{Z}^d \mid \tilde{\tau}_i^\ell + 1 \leq \tau_i(\beta) \leq \tilde{\tau}_i^{\ell+1}\}$$

for $\ell = 0, \dots, r$. At most $r + 1$ of these strips are nonempty, because some of the hyperplanes may coincide. Also, the last few of the $\tilde{\tau}_i^\ell$ will equal ∞ ; we interpret any strip where $\tau_i^\ell = \tau_i^{\ell+1} = \infty$ as empty, and ignore it.

PROPOSITION 6.2. *Let J be as in (6.2). For any fixed $\ell_1, \dots, \ell_n \in \{0, \dots, r\}$, the lattice points in the polyhedron*

$$\Delta(\ell_1, \dots, \ell_n) := \bigcap_{i=1}^n \{\beta \in \mathbb{R}^d \mid \tilde{\tau}_i^{\ell_i} + 1 \leq \tau_i(\beta) \leq \tilde{\tau}_i^{\ell_i+1}\}$$

all lie inside a single sector in $\mathcal{S}(J)$. The partition of \mathbb{Z}^d by the polyhedra $\Delta(\ell_1, \dots, \ell_n)$ refines the partition of \mathbb{Z}^d by the sectors in $\mathcal{S}(J)$.

PROOF. This follows from the definitions and (6.1), which uses that Q is saturated. \square

Proposition 6.2 makes way for an algorithm to compute the set of sectors.

ALGORITHM 6.3.

```

INPUT    $J = \bigoplus_{j=1}^r J^j$ , an injective module over  $k[Q]$ , with  $J^j = k\{\alpha_j + F_j - Q\}$ 
OUTPUT  the set  $\mathcal{S}(J)$  of sectors, each expressed as a list of polyhedra that partition it
DEFINE   $\phi: \mathbb{Z}^d \rightarrow$  subsets of  $\{1, \dots, r\}$ , as in Lemma 6.1
INITIALIZE  $\mathcal{A} := \{\}$ , the empty collection of subsets of  $\{1, \dots, r\}$ 
WHILE   $\ell_1, \dots, \ell_n \in \{0, \dots, r\}$  DO
  IF  $\Delta(\ell_1, \dots, \ell_n) \neq \emptyset$ 
    THEN DEFINE  $A := \phi(\Delta(\ell_1, \dots, \ell_n))$ 
    ELSE NEXT  $(\ell_1, \dots, \ell_n)$ 
  END IF-THEN-ELSE
  IF  $A \in \mathcal{A}$ 
    THEN REDEFINE  $S_A := S_A \cup \{\Delta(\ell_1, \dots, \ell_n)\}$ 
    ELSE INITIALIZE  $S_A := \{\Delta(\ell_1, \dots, \ell_n)\}$ 
    REDEFINE  $\mathcal{A} := \mathcal{A} \cup \{A\}$ 
  END IF-THEN-ELSE
NEXT  $(\ell_1, \dots, \ell_n)$ 
END WHILE-DO
OUTPUT  $\{S_A \mid A \in \mathcal{A}\}$ 

```

Note that ϕ is constant on $\Delta(\ell_1, \dots, \ell_n)$ by definition, and can easily be determined directly from the data (ℓ_1, \dots, ℓ_n) .

Next comes the determination of which maps \mathbf{x}^{B-A} need computing. In the coming algorithm, we write $\Delta(\ell_1, \dots, \ell_n) \leq \Delta(\ell'_1, \dots, \ell'_n)$ if $(\ell_1, \dots, \ell_n) \leq (\ell'_1, \dots, \ell'_n)$ as vectors in $(\mathbb{Z} \cup \infty)^n$. Such notation is justified because $\Delta(\ell_1, \dots, \ell_n) \not\leq \Delta(\ell'_1, \dots, \ell'_n)$ automatically implies that $\Delta(\ell'_1, \dots, \ell'_n) - \Delta(\ell_1, \dots, \ell_n)$ fails to intersect Q .

ALGORITHM 6.4.

```

INPUT  sectors  $S_A$  and  $S_B$  in  $\mathcal{S}(J)$  from the output of Algorithm 6.3
OUTPUT the truth value of: “there exist  $\alpha \in S_A$  and  $\beta \in S_B$  with  $\beta - \alpha \in Q$ ”
INITIALIZE  $val := \text{FALSE}$ 
WHILE  $(\Delta_A, \Delta_B) \in S_A \times S_B$  AND  $val = \text{FALSE}$ , DO
  IF  $A \supseteq B$  AND  $\Delta_A \leq \Delta_B$ 
    THEN DEFINE  $\Delta_B - \Delta_A :=$  the Minkowski sum of  $\Delta_B$  and  $-\Delta_A$ 
    ELSE NEXT  $(\Delta_A, \Delta_B)$ 
  END IF-THEN-ELSE
  IF  $Q \cap (\Delta_B - \Delta_A) \neq \emptyset$ 
    THEN REDEFINE  $val := \text{TRUE}$ 
    ELSE NEXT  $(\Delta_A, \Delta_B)$ 
  END IF-THEN-ELSE
END WHILE-DO
OUTPUT  $val$ 

```

The proof of correctness for Algorithm 6.4 is straightforward from the definitions, except for the first IF-THEN-ELSE procedure, which relies on Lemma 6.5, below. Note the non-necessity in Algorithm 6.4 of actually finding a witness in $\Delta_B - \Delta_A$ for $S_A \preceq S_B$; as we have seen in (5.1) and (5.2) from the proof of Theorem 5.2, the natural map on cohomology is induced by taking submatrices of the monomial matrix, regardless of where the witnesses lie.

LEMMA 6.5. *If $\mathcal{S}(J)$ is as in Proposition 5.1, then $Q \cap (S_B - S_A) \neq \emptyset$ implies $A \supseteq B$.*

PROOF. If $a \in Q$ and $(J_i)_\alpha = 0$, then $(J_i)_{a+\alpha} = 0$, so the set of summands nonzero in degree $a + \alpha$ can only be smaller. \square

Unfortunately, Algorithm 6.4 is necessary, because $\Delta_B - \Delta_A \neq \emptyset$ need not always hold when $\Delta_A \leq \Delta_B$, as the example to come shortly demonstrates. It does seem, however, that the offending pairs of polytopes are usually “small”. For instance, we know of no examples where the lattice points in either polytope affinely span \mathbb{Z}^d .

EXAMPLE 6.6. Let Q be the subsemigroup of \mathbb{N}^2 generated by $(2, 0)$, $(1, 1)$, and $(0, 2)$. Name the faces of Q as $\mathbf{0}, X, Y, Q$, and set $E_F = F - Q$. Let

$$\begin{aligned}
 J = & \mathbf{k}\{(0, 0) + E_{\mathbf{0}}\} \oplus \mathbf{k}\{(0, 1) + E_X\} \oplus \mathbf{k}\{(0, 0) + E_Y\} \\
 & \oplus \mathbf{k}\{(0, -1) + E_X\} \oplus \mathbf{k}\{(-2, 0) + E_Y\},
 \end{aligned}$$

with the summands labeled in order as J_1, \dots, J_5 . Letting X be facet number 1 and Y be facet number 2, the arrays τ_i^j and $\tilde{\tau}_i^\ell$ look like

$$\begin{pmatrix} \tau_1^j \\ \tau_2^j \end{pmatrix} = \begin{pmatrix} 0 & 1 & \infty & -1 & \infty \\ 0 & \infty & 0 & \infty & -2 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \tilde{\tau}_1^\ell \\ \tilde{\tau}_2^\ell \end{pmatrix} = \begin{pmatrix} -\infty & -1 & 0 & 1 & \infty & \infty & \infty \\ -\infty & -2 & 0 & 0 & \infty & \infty & \infty \end{pmatrix}$$

The sectors $S_{\{1,2,3\}}$ and $S_{\{2,3\}}$ contain one polytope each, and both of these polytopes contain exactly one lattice point. Specifically, identifying the sector, the polytope, and the lattice point, we have

$$S_{\{1,2,3\}} = \Delta(-1, -2) = (0, 0) \quad \text{and} \quad S_{\{2,3\}} = \Delta(0, -2) = (-1, 1).$$

Now $\Delta(-1, -2) \leq \Delta(0, -2)$, but subtracting the vector in $S_{\{1,2,3\}}$ from the one in $S_{\{2,3\}}$ yields $(-1, 1)$, which does not lie in the semigroup Q .

REMARK 6.7. The notion of sector partition ought to have a refinement that takes into account the various kinds of failures of saturation for arbitrary affine semigroup. The resulting notion would produce sector partitions for the cohomology of complexes of injectives over nonnormal affine semigroup rings. The failures of saturation fall into two categories: the geometric kind, arising from polyhedral “holes” in the semigroup (as compared with its saturation), and the arithmetic kind, arising from finite-index sublattices generated by faces. Even in the case where arithmetic failure is absent, however, we do not know how to bound the sizes and shapes of the “holes” sufficiently to carry out an analysis such as the one producing the algorithms above.

7. Computing local cohomology with monomial support

Still assuming that Q is sturated, we have now finally developed enough tools to prove the main theorem on local cohomology with monomial support, namely Theorem 1.3 from the Introduction.

PROOF OF THEOREM 1.3. Take $i = d$ in Theorem 4.7, and let $J^\bullet(-a)$ be the minimal \mathbb{Z}^d -graded injective resolution computed there. Then J^\bullet is an algorithmically computed injective resolution of M . By definition, $H_I^i(M)$ is the middle cohomology of the complex $\Gamma_I J^{i-1} \rightarrow \Gamma_I J^i \rightarrow \Gamma_I J^{i+1}$, where $\Gamma_I J^j$ is the direct sum of all indecomposable summands of J^j whose unique associated prime contains I . Having now expressed $H_I^i(M)$ as the cohomology of an effectively computed complex of injectives, Theorem 5.2 says that $H_I^i(M)$ has a sector partition. The set of sectors in part 1 of Definition 1.2 is computed by Algorithm 6.3. The vector spaces in part 2 of Definition 1.2 are specified in (5.1) from the proof of Theorem 5.2, and naturally determine the maps in part 3 of Definition 1.2, given the computation in Algorithm 6.4. \square

Now we turn to issues of complexity. There is little sense in completing a formal complexity analysis of all of the algorithms presented in this paper, as they involve Gröbner basis computation, which is doubly-exponential from a worst-case perspective. However, it is worth mentioning where the complexity in our algorithms comes from, up to a factor arising from the complexity of Gröbner basis computation, since Gröbner basis computations are often more efficient than expected. The purpose of what follows, therefore, is to assure the reader that our algorithms have not amplified the faux-doubly-exponential complexity of Gröbner bases with some “honest” exponential complexity.

Let us assume that the dimension d is fixed, and analyze the complexity of computing all of the local cohomology of a finitely generated module M supported on a fixed monomial ideal I over a normal semigroup ring $k[Q]$. This computation involves all of the algorithms in the paper except the one in Section 3.3. (The complexity of Algorithm 3.15 above and beyond Algorithm 3.6 is only about as bad as that of $k[Q^{\text{sat}}]/k[Q]$ as a $k[Q]$ -module, anyway.)

In Algorithm 3.6, the only non-Gröbner contribution to the running time comes from the number of basis elements constructed (see Remark 3.8.2, which can be used to ensure that we only check faces of Q giving rise to basis elements). This number is by definition

a Bass number of M . Thus, up to Gröbner basis computation, Algorithm 3.6 is only as complex as its output.

Next we consider the algorithm in Proposition 3.14. The algorithm works by taking the union (over a set of facets of Q) of ideals output by Algorithm 3.11. The output presents the generators of each such ideal as the lattice points in a union of polytopes having the form $((a + \mathbb{R}H) \cap \mathbb{R}_+D) + G$, where D is a face of Q . The computation of each such polytope is by standard techniques to intersect polyhedra and take Minkowski sums with the fixed zonotope G . Hence, up to factors coming from the number of facets of Q and from standard procedures, we need only bound

1. the number of polytopes output by Algorithm 3.11, and
2. the number of lattice points in each such polytope.

The former is polynomial in the number of facets of Q by Remark 3.13.2. The latter is polynomial in the input vector $a \in Q$ by the piecewise polynomiality of the lattice point enumeration function of $(a + \mathbb{R}H) \cap \mathbb{R}_+D$ as a function of a (McM77), along with the fact that G is fixed. Actually computing the set of lattice points in each polytope can be accomplished using the efficient algorithms of Barvinok and Woods (BW02).

REMARK 7.1. We need to do Gröbner basis computations with the irreducible ideals W output by the algorithm in Proposition 3.14. This means that, for our purposes, the *short rational generating functions* output by the algorithms of (BW02) do not suffice: we actually require the list of lattice point explicitly, to get a generating set of W as a list of monomials. Thus the short generating functions must be expanded. To reduce complexity, the short generating functions can be post-processed using the methods of (BW02) to yield short generating functions for the *minimal* generators of the ideals W in question. Then we can expand only these “minimal” short generating functions.

The remaining contributions to the complexity of our local cohomology computation come from Algorithm 6.3, which computes the sets of polytopes whose disjoint unions constitute the sectors, and Algorithm 6.4. The latter is quadratic in the output of Algorithm 6.3, times a factor coming from the Minkowski sum operations and the decision procedure for whether each such sum contains a lattice point after intersecting with Q . Therefore it remains only to analyze Algorithm 6.3.

PROPOSITION 7.2. *The number of polyhedra arising in Algorithm 6.3 is polynomial in the Bass numbers of M and the number of facets of Q .*

PROOF. Each Bass number of M represents an indecomposable injective module whose bounding hyperplanes subdivide \mathbb{R}^d into a number of regions. Consider the subdivision of \mathbb{R}^d obtained by taking simultaneously all of the hyperplanes corresponding all of the Bass numbers of M . The number of hyperplanes contributed by each Bass number is at most the number of facets of Q , so the total number of hyperplanes is at most the number of facets of Q times the sum of the contributing Bass numbers. It is well known (and follows by induction on n and the dimension d) that n hyperplanes subdivide \mathbb{R}^d into a number of regions that is a polynomial in n of degree d . \square

This proof shows that the number of polyhedra is exponential in the dimension. Exponential growth as a function of dimension also occurs in the analysis before Remark 7.1, where we apply (McM77).

REMARK 7.3. A large number of rational polyhedra arise in the course of computing local cohomology modules. When the identification of all the lattice points in these polyhedra is necessary, the complexity of this task should be drastically reduced by the fact that most of these polyhedra have facets parallel to those of Q itself. Results such as those in (BV97) could be helpful along these lines.

ACKNOWLEDGMENTS

We wish to thank Gennady Lyubeznik, who motivated us to make our previously abstract methods algorithmic. Both authors were partially supported by the National Science Foundation.

References

- Winfried Bruns and Jürgen Herzog, *Cohen–Macaulay rings*, Cambridge Studies in Advanced Mathematics, vol. 39, Cambridge University Press, Cambridge, 1993.
- Michel Brion and Michèle Vergne, *Residue formulae, vector partition functions and lattice points in rational polytopes*, J. Amer. Math. Soc. **10** (1997), no. 4, 797–833.
- Alexander Barvinok and Kevin Woods, *Short rational generating functions for lattice point problems*, J. Amer. Math. Soc., to appear. [arXiv:math.CO/0211146](https://arxiv.org/abs/math/0211146)
- David Eisenbud, *Commutative algebra, with a view toward algebraic geometry*, Graduate Texts in Mathematics, vol. 150, Springer-Verlag, New York, 1995, first printing.
- David Eisenbud, Mircea Mustață, and Mike Stillman, *Cohomology on toric varieties and local cohomology with monomial supports*, J. Symbolic Comput. **29** (2000), no. 4-5, 583–600.
- Robin Hartshorne, *Affine duality and cofiniteness*, Invent. Math. **9** (1969/1970), 145–164.
- David Helm and Ezra Miller, *Bass numbers of semigroup-graded local cohomology*, Pacific J. Math. **209** (2003), no. 1, 41–66.
- Craig Huneke and Rodney Sharp, *Bass numbers of local cohomology modules*, Trans. Amer. Math. Soc. **339** (1993), 765–779.
- Gennady Lyubeznik, *Finiteness properties of local cohomology modules (an application of D -modules to commutative algebra)*, Invent. Math. **113** (1993), no. 1, 41–55.
- P. McMullen, *Valuations and Euler-type relations on certain classes of convex polytopes*, Proc. London Math. Soc. (3) **35** (1977), no. 1, 113–135.
- Ezra Miller, *Alexander duality for monomial ideals and their resolutions*, [math.AG/9812095](https://arxiv.org/abs/math/9812095), 1998.
- Ezra Miller, *The Alexander duality functors and local duality with monomial support*, J. Algebra **231** (2000), 180–234.
- Ezra Miller, *Cohen–Macaulay quotients of normal semigroup rings via irreducible resolutions*, Math. Res. Lett. **9** (2002), no. 1, 117–128.
- Ezra Miller and Bernd Sturmfels, *Combinatorial commutative algebra*, version of 30 July 2003 (currently available at <http://math.umn.edu/~ezra>), in progress.
- Mircea Mustață, *Local cohomology at monomial ideals*, J. Symbolic Comput. **29** (2000), no. 4-5, 709–720.
- Naoki Terai, *Local cohomology modules with respect to monomial ideals*, preprint, 1999.
- Wolmer V. Vasconcelos, *Computational methods in commutative algebra and algebraic geometry*, Algorithms and Computation in Mathematics, vol. 2, Springer-Verlag, Berlin, 1998, With chapters by David Eisenbud, Daniel R. Grayson, Jürgen Herzog and Michael Stillman.
- Uli Walther, *Algorithmic computation of local cohomology modules and the cohomological dimension of algebraic varieties*, J. Pure Appl. Algebra **139** (1999), 303–321.
- Kohji Yanagawa, *Sheaves on finite posets and modules over normal semigroup rings*, J. Pure Appl. Algebra **161** (2001), no. 3, 341–366.
- Kohji Yanagawa, *Squarefree modules and local cohomology modules at monomial ideals*, Local cohomology and its applications (Guanajuato, 1999), Lecture Notes in Pure and Appl. Math., vol. 226, Dekker, New York, 2002, pp. 207–231.