

Computing Multidimensional Persistence*

Gunnar Carlsson
Dept. of Mathematics
Stanford University

Gurjeet Singh
Dept. of Mathematics
Stanford University

Afra Zomorodian
Dept. of Computer Science
Dartmouth College

Preliminary Manuscript – 10/2/09

Abstract

The theory of multidimensional persistence captures the topology of a multifiltration – a multiparameter family of increasing spaces. Multifiltrations arise naturally in the topological analysis of scientific data. In this paper, we give a polynomial time algorithm for computing multidimensional persistence. We recast this computation as a problem within computational commutative algebra and utilize algorithms from this area to solve it. While the resulting problem is EXSPACE-complete and the standard algorithms take doubly-exponential time, we exploit the structure inherent withing multifiltrations to yield practical algorithms. We implement all algorithms in the paper and provide statistical experiments to demonstrate their feasibility.

1 Introduction

In this paper, we give a polynomial time algorithm for computing the persistent homology of a multifiltration. The computed solution is compact and complete, but not an invariant. Theoretically, this is the best one may hope for because complete compact invariants do not exist for multidimensional persistence. We also discuss computing an incomplete invariant, the rank invariant, giving an algorithm for reading off this invariant from our complete solution. We implement all algorithms in the paper and provide statistical experiments to demonstrate their feasibility.

1.1 Motivation

Intuitively, a multifiltration models a growing space that is parameterized along multiple dimensions. For example, the complex with coordinate $(3, 2)$ in Figure 1 is filtered along the horizontal and vertical dimensions, giving rise to a bifiltration. Multifiltrations arise naturally in topological analysis of scientific data. Often, scientific data is in the form of a finite set of noisy samples from some underlying topological space. Our goal is to robustly recover the lost connectivity of the underlying space. If the sampling is dense enough, we approximate the space as a union of balls by placing ϵ -balls around each point. As we increase ϵ , we obtain a growing family of spaces, a one-parameter multifiltration also called a filtration. This approximation is the central idea behind many methods for computing the topology of a point set, such as *Čech*, *Rips-Vietoris* [12], or *witness* [7] complexes. Often, however, the input point set is filtered via multiple functions. For example, in analyzing the structure of natural images, we filter the data according to density [1]. We now have multiple dimensions along which our space is filtered; that is, we have a multifiltration.

1.2 Prior Work

For one-dimensional filtrations, the theory of persistent homology provides a complete invariant called a *barcode*, a multiset of intervals [22]. Each interval in the barcode corresponds to the lifetime of a single topological feature within the filtration. Since intrinsic features have long lives, while noise is short-lived, a quick examination of the intervals gives a robust estimation of the topology. The existence of a complete compact invariant, as well as efficient algorithms and fast implementations have led to successful application of persistence to a variety of disciplines, such

*The authors were partially supported by the following grants: G. C. by NSF DMS-0354543; A. Z. by DARPA HR 0011-06-1-0038, ONR N 00014-08-1-0908, and NSF CCF-0845716; all by DARPA HR 0011-05-1-0007.




Figure 1: A bifiltration. The complex with labeled vertices is at coordinate $(3, 2)$. Simplices are highlighted and named at the critical coordinates that they appear.

as shape description [5], denoising volumetric density data [13], detecting holes in sensor networks [8], analyzing neural activity in the visual cortex [18], and analyzing the structure of natural images [1], to name a few.

For multifiltrations of dimension higher than one, the situation is much more complicated. The theory of *multidimensional persistence* shows that no complete discrete invariant exists, where *discrete* means that the structure of the target for the invariant does not depend on the structure of the underlying field of coefficients [2]. Instead, the authors propose an incomplete invariant, the rank invariant, which captures important persistent information. Unfortunately, this invariant is not compact, requiring large storage, so its direct computation using the one-dimensional algorithm is not feasible. A variant of the problem of multidimensional persistence appears in computer vision [10]. There is also a partial solution called *vineyards* [4]. A full solution, however, has not been attempted by any prior work.

1.3 Contributions

In this paper, we provide a complete solution to the problem of computing multidimensional persistence.

- We recast persistence as a problem within computational commutative algebra, allowing us to utilize powerful algorithms from this area.
- We exploit the structure provided by a multifiltration to greatly simplify the algorithms.
- We show that the resulting algorithms are polynomial time, unlike their original counterparts, which are EXSPACE-complete, requiring exponential space and time.
- We implement all algorithms and show that the multidimensional setting requires different data structures than the one-dimensional case for efficient computation. In particular, the change in approach allows for parallelization.
- We provide a suite of statistical tests with random multifiltrations and analyze the running time in practice.

As such, our work provides a full integrated solution, from a new theoretical understanding in computational commutative algebra, to the refinement of algorithms, full implementation, and experimentation.

We begin with background on multidimensional persistence in Section 2. In Section 3, we formalize the input to our algorithms and justify it. In Section 4, we reinterpret the problem of computing multidimensional persistence within computational commutative algebra. Having recast the problem, we use algorithms from this area to solve the problem. This gives us a computationally intractable solution. In Section 5, we simplify the algorithms by using the structure within multifiltrations. This simplification allows us to derive a polynomial time algorithm from the original doubly-exponential algorithms. In Section 7, we describe our implementations of these algorithms and provide experiments to validate our work in practice.

2 Background

In this section, we provide the background on multidimensional persistence. We begin by formalizing multifiltrations. We then describe simplicial homology, persistent homology, and multidimensional persistence, as well as the definition of the rank invariant.

2.1 Multifiltrations

Let $\mathbb{N} \subseteq \mathbb{Z}$ be the set of non-negative integers. For vectors in \mathbb{N}^n or \mathbb{R}^n , we say $u \leq v$ if $u_i \leq v_i$ for $1 \leq i \leq n$. A topological space X is *multifiltered* if we are given a family of subspaces $\{X_u\}_u$, where $u \in \mathbb{N}^n$ and $X_u \subseteq X$ such that for $u, w_1, w_2, v \in \mathbb{N}^n$, the diagrams

$$\begin{array}{ccc} X_u & \longrightarrow & X_{w_1} \\ \downarrow & & \downarrow \\ X_{w_2} & \longrightarrow & X_v \end{array} \quad (1)$$

commute whenever $u \leq w_1, w_2 \leq v$. We call the family of subspaces $\{X_u\}_u$ a *multifiltration*. A one-dimensional multifiltration is called a *filtration*. We show an example of a bifiltration of a simplicial complex in Figure 1. In a multifiltration, any path with monotonically increasing coordinates is a filtration, such as any row or column in the figure. Multifiltrations constitute the input to our algorithms. We further motivate their use as a model for scientific data in the next section.

2.2 Homology

Given a topological space, *homology* is a topological invariant that is often used in practice as it is easily computable. Here, we describe *simplicial homology* briefly, referring the reader to Hatcher [14] as a resource. We assume our input is a simplicial complex K , such as the complexes in Figure 1. We note, however, that our results carry over to arbitrary cell complexes, so we use cell complexes when appropriate in our presentation. The *i*th chain group $C_i(K)$ of K is the free Abelian group on K 's set of oriented *i*-simplices. An element $c \in C_i(K)$ is an *i*-chain, $c = \sum_j n_j [\sigma_j]$, $\sigma_j \in K$ with *coefficients* $n_j \in \mathbb{Z}$. Given such a chain c , the *boundary operator* $\partial_i: C_i(K) \rightarrow C_{i-1}(K)$ is a homomorphism defined linearly by its action on any simplex $\sigma = [v_0, v_1, \dots, v_i] \in c$,

$$\partial_i \sigma = \sum_j (-1)^j [v_0, \dots, \hat{v}_j, \dots, v_i], \quad (2)$$

where \hat{v}_j indicates that v_j is deleted from the vertex sequence. The boundary operator connects the chain groups into a *chain complex* C_* :

$$\dots \rightarrow C_{i+1}(K) \xrightarrow{\partial_{i+1}} C_i(K) \xrightarrow{\partial_i} C_{i-1}(K) \rightarrow \dots \quad (3)$$

Using the boundary operator, we may define subgroups of C_i : the *cycle group* $\ker \partial_i$ and the *boundary group* $\text{im } \partial_{i+1}$. Since $\partial_i \partial_{i+1} = 0$, then $\text{im } \partial_{i+1} \subseteq \ker \partial_i \subseteq C_i(K)$. The *i*th homology group is

$$H_i(K) = \ker \partial_i / \text{im } \partial_{i+1} \quad (4)$$

and the *i*th Betti number is $\beta_i(K) = \text{rank } H_i(K)$. Over field coefficients k , H_i is a k -vector space of dimension β_i .

2.3 Persistent Homology

Given a multifiltration $\{X_u\}_u$, the homology of each subspace X_u over a field k is a vector space. For instance, the bifiltered complex in Figure 1 has zeroth homology vector spaces isomorphic to the commutative diagram

$$\begin{array}{ccccccc}
 k^2 & \longrightarrow & k^2 & \longrightarrow & k & \longrightarrow & k \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 k^2 & \longrightarrow & k^4 & \longrightarrow & k^3 & \longrightarrow & k^3 \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 k^4 & \longrightarrow & k^3 & \longrightarrow & k^2 & \longrightarrow & k^2
 \end{array}$$

where the dimension of the vector space counts the number of components of the complex, and the maps between the homology vector spaces are induced by the inclusion maps relating the subspaces. Persistent homology captures information contained in the induced maps. There are two equivalent definitions that we use in this paper. The first definition was originally for filtrations only [9], but was later extended to multifiltrations [2]. The key idea is to relate homology between a pair of complexes. For each pair $u, v \in \mathbb{N}^n$ with $u \leq v$, $X_u \subseteq X_v$ by definition, so $X_u \hookrightarrow X_v$. This inclusion, in turn, induces a linear map $\iota_i(u, v)$ at the i th homology level $H_i(X_u) \rightarrow H_i(X_v)$ that maps a homology class within X_u to the one that contains it within X_v . The i th persistent homology is $\text{im } \iota_i$, the image of ι_i for all pairs $u \leq v$. This definition also enables the definition of an incomplete invariant. The i th rank invariant is

$$\rho_i(u, v) = \text{rank } \iota_i(u, v), \quad (5)$$

for all pairs $u \leq v \in \mathbb{N}^n$, where ι_i is defined above [2].

While the above definition provides intuition, it is inexpedient for theoretical development. For most of our paper, we use a second definition of persistence that is grounded in algebraic topology, allowing us to utilize tools from commutative algebra for computation [22, 2].

2.4 Multidimensional Persistence

The key insight for the second definition is that the persistent homology of a multifiltration is, in fact, the homology of a single algebraic entity. This object encodes all the complexes via polynomials that track cells through the multifiltration. To define our algebraic structure, we need to first review graded modules over polynomials. A *monomial* in x_1, \dots, x_n is a product of the form

$$x_1^{v_1} \cdot x_2^{v_2} \cdot \dots \cdot x_n^{v_n} \quad (6)$$

with $v_i \in \mathbb{N}$. We denote it x^v , where $v = (v_1, \dots, v_n) \in \mathbb{N}^n$. A *polynomial* f in x_1, \dots, x_n and *coefficients* in field k is a finite linear combination of monomials, $f = \sum_v c_v x^v$, with $c_v \in k$. The set of all such polynomials is denoted $k[x_1, \dots, x_n]$. For instance, $5x_1x_2^2 - 7x_3^3 \in k[x_1, x_2, x_3]$. An n -graded ring is a ring R equipped with a decomposition of Abelian groups $R \cong \bigoplus_v R_v$, $v \in \mathbb{N}^n$ so that multiplication has the property $R_u \cdot R_v \subseteq R_{u+v}$. Elements in a single group R_u are called *homogeneous*. The set of polynomials form the n -graded *polynomial ring*, denoted A^n . This ring is graded by $A_v^n = kx^v$, $v \in \mathbb{N}^n$. An n -graded module over an n -graded ring R is an Abelian group M equipped with a decomposition $M \cong \bigoplus_v M_v$, $v \in \mathbb{N}^n$ together with a R -module structure so that $R_u \cdot M_v \subseteq M_{u+v}$. An n -graded module is *finitely generated* if it admits a finite generating set. Also, recall the notion of a *free* module on an n -graded set and a *basis* for such a module [2].

Given a multifiltration $\{X_u\}_u$, the i th dimensional homology is the following n -graded module over A^n

$$\bigoplus_u H_i(X_u), \quad (7)$$

where the k -module structure is the direct sum structure and $x^{v-u}: H_i(X_u) \rightarrow H_i(X_v)$ is the induced homomorphism $\iota_i(u, v)$ we described in the previous section. This view of homology yields two important results. In one dimension, the persistent homology of a filtration is easily classified and parameterized by the barcode, and there is an efficient algorithm for its computation [22]. In higher dimensions, no similar classification exists [2]. Instead, we may utilize an incomplete invariant. For instance, the rank invariant defined above is provably equivalent to the barcode in one dimension, allowing us to compute incomplete information in higher dimensions.

3 One-Critical Multifiltrations

Our interest in persistent homology stems from the need for analyzing scientific data. In this section, we begin by formalizing such data. We then show that topological analysis of scientific data naturally generates multifiltrations. In particular, it generates multifiltrations with the following property.

Definition 1 (one-critical) *A multifiltered complex K where each cell σ has a unique critical coordinate u_σ is one-critical.*

In the rest of this paper, we assume that our input multifiltrations are one-critical. General multifiltrations, however, may not have this property. Therefore, we end this section by describing a classic construction that eliminates multiple critical coordinates in such input.

3.1 Model for Scientific Data

We are often given scientific data in the form of a set of samples, possibly noisy, from some underlying geometric space. At each sample point, we may also have measurements from the ambient space. For example, a fundamental goal in graphics is to render objects under different lighting from different camera positions. One approach is to construct a digitized model using data from a range scanner, which employs multiple cameras to sense 3D positions on an object's surface, as well as estimated normals and texture information [19]. An alternate approach samples the four-dimensional *light field* of a surface directly and interpolates to render the object without explicit surface reconstruction [15]. Either approach gives us a set of noisy samples with measurements. Similarly, a *node* in a *wireless sensor network* has sensors on board that measure physical attributes of the local environment, such as pressure and temperature [21]. The GPS coordinates of the nodes give us a set of samples at which several functions are sampled.

Formally, we assume we have a manifold X with $d - 1$ Morse functions defined on it [16]. In practice, X is often embedded within a high-dimensional Euclidean space \mathbb{R}^d , although this is not required. As such, we model the data using the following definition.

Definition 2 (multifiltered data) *We assume our input is multifiltered data, a finite set S of d -dimensional points along with $n - 1$ real-valued functions $f_j: S \rightarrow \mathbb{R}$, $1 \leq j \leq n - 1$, defined at the samples, for $n > 1$.*

3.2 Construction

We begin by approximating the underlying space X with a combinatorial representation, a *complex*. There are a variety of techniques for building such complexes, all of which have a *scale* parameter ϵ , such as the *Čech* and *Rips-Vietoris* [12] complex which employ a global parameter, or the *witness complex* [7] which simulates a local parameter. As we increase ϵ , a complex grows larger, and fixing a maximum scale ϵ_{\max} gives us a filtered complex K with n cells. Each cell $\sigma \in K$ enters K at scale $\epsilon(\sigma)$. We formalize this type of complex next.

Definition 3 (scale-filtered complex) *A scale-filtered complex is the tuple (K, σ) , where K is a finite complex and $\epsilon: K \rightarrow \mathbb{R}$, such that for any $\sigma \in K$, $\sigma \in K_u$ only for $u \geq \epsilon(\sigma)$.*

We now assume we have a scale-filtered complex K, ϵ defined on our multifiltered point set. To incorporate the measurements at the points into our analysis, we first extend the sampled functions f_j to the cells in the complex. For $\sigma \in K$ and f_j , let $f_j(\sigma)$ be the maximum value f_j takes on σ 's vertices; that is, $f_j(\sigma) = \max_{v \in \sigma} f_j(v)$, where v is a vertex. This extension defines $n - 1$ functions on the complex, $f_j: K \rightarrow \mathbb{R}$. We now combine all filtration functions into a single multivariate function $F: K \rightarrow \mathbb{R}^n$, where

$$F(\sigma) = (f_1(\sigma), f_2(\sigma), \dots, f_{n-1}(\sigma), \epsilon(\sigma)). \quad (8)$$

We may now filter K via the *excursion sets* $\{K_u\}_u$ of F :

$$K_u = \{\sigma \in K \mid F(\sigma) \leq u \in \mathbb{R}^n\}. \quad (9)$$

Note that simplex σ enters K_u at $u = F(\sigma)$ and will remain in the complex for all $u \geq F(\sigma)$. Equivalently, $F(\sigma)$ is the unique minimal *critical* coordinate at which σ enters the filtered complex. That is, these multifiltrations are one-critical.

Example 1 (bifiltration criticals) The bifiltration in Figure 1 is one-critical, as each simplex enters at a single critical coordinate. For instance, $F(a) = (1, 1)$, $F(cde) = (3, 1)$, and $F(af) = (1, 2)$.

Since K is finite, we have a finite set of critical coordinates that we may project on each dimension j , getting a finite set of critical values C_j . We now restrict ourselves to the Cartesian product $C_1 \times \dots \times C_n$ of the critical values, parameterizing the resulting discrete grid using \mathbb{N} in each dimension. It is clear that the required diagrams (1) commute, giving us a d -dimensional multifiltration $\{K_v\}_v$ with $v \in \mathbb{N}^n$.

We end by noting that one-critical multifiltrations may be represented compactly by the set of tuples

$$\{(\sigma, F(\sigma)) \mid \sigma \in K\}. \quad (10)$$

This representation is the main input to our algorithms in Section 4.3.

3.3 Mapping Telescope

In general, multifiltrations are not one-critical since a cell may have enter at multiple incomparable critical coordinates, viewing \leq as a partial order on \mathbb{N}^n . For example, in Figure 1, the vertex that enters at $(1, 0)$ may also enter at $(0, 1)$ as the two coordinates are incomparable. For such multifiltrations, we may utilize the *mapping telescope*, a standard algebraic construction to ensure that each cell has a unique critical coordinate [14]. Intuitively, this construction introduces additional *shadow* cells into the multifiltration without changing its topology. We will not detail this construction here as none of the multifiltrations we encounter in practice require the conversion. We should note, however, that the mapping telescope increases the size of the multifiltration, depending on the number of cells with multiple critical points. In the worst case, the increase is exponential, but this requires strange constructions that are not topologically interesting and do not arise in practice.

4 Using Computational Commutative Algebra

Having described our input in the last section, in this section we recast the problem of computing multidimensional persistence as a problem within computational commutative algebra. We then describe standard algorithms from this area that solve our problem. While this gives us a solution, it comes at a computational cost, as these algorithms are intractable. In the next section, we refine them to derive polynomial-time algorithms.

4.1 Multigraded Homology

We begin by extending simplicial homology to a multifiltered simplicial complex. We then convert the computation of the latter to standard questions in computational commutative algebra.

Definition 4 (chain module) *Given a multifiltered simplicial complex $\{K_u\}_u$, the i th chain module is the n -graded module over A^n*

$$C_i = \bigoplus_u C_i(K_u), \quad (11)$$

where the k -module structure is the direct sum structure and $x^{v-u}: C_i(K_u) \rightarrow C_i(K_v)$ is the inclusion $K_u \hookrightarrow K_v$.

Note that we overload notation to reduce complexity by having $C_i = C_i(\{K_u\}_u)$ when the multifiltration is clear from context. The module C_i is n -graded as for any $u \in \mathbb{N}^n$, $(C_i)_u = C_i(K_u)$. That is, the chain complex in grade u of the module is the chain complex of K_u , the simplicial complex with coordinate u .

Example 2 (bifiltration module) Consider the vertex d in the bifiltered complex in Figure 1. This vertex has critical coordinate $(1, 0)$, so copies of this vertex exist in 9 complexes K_u for $u \geq (1, 0)$. The inclusion maps relate these copies within the complexes. In turn, polynomials relate the chain groups in the different grades of the module. Let d be the copy of the vertex in the critical coordinate $(1, 0)$. Then, within C_i , we have d in grade $(1, 0)$, $x_1 d$ in grade $(2, 0)$, $x_2 d$ in grade $(1, 1)$, $x_1^2 x_2^2 d$ in grade $(3, 2)$ and so on, as required by the definition of an n -graded module. In other words, a simplex has different names in different grades.

The graded chain modules C_i are finitely generated, so we may choose bases for them.

Definition 5 (standard basis) The standard basis for the i^{th} chain module C_i is the set of i -simplices in critical grades.

Example 3 (bifiltration bases) For our bifiltration in Figure 1, the standard basis corresponds to the highlighted and named simplices. For example, the standard basis for C_0 is

grade	(0, 0)	(1, 0)	(1, 1)
simplices	b, c, e, f	d	a

Note that in doing so, we have made a choice of basis. Unlike for chain groups, this choice has an important consequence: Our resulting calculations will not be invariant but depend on the initial basis.

Recall that our multifiltrations are one-critical. The graded chain groups of one-critical multifiltrations are free: Since each cell enters only once, the resulting chain groups do not require any relations. Since our graded chain groups are free, the boundary operator is simply a homomorphism between free graded modules. Given standard bases, we may write the boundary operator $\partial_i : C_i \rightarrow C_{i-1}$ explicitly as a matrix with polynomial entries.

Example 4 (boundary matrix) For the bifiltration in Figure 1, ∂_1 has the matrix

$$\begin{bmatrix} & ab & bc & cd & de & ef & af & bf & ce \\ a & x_2 & 0 & 0 & 0 & 0 & x_2 & 0 & 0 \\ b & x_1x_2^2 & x_1^2x_2^2 & 0 & 0 & 0 & 0 & x_2^2 & 0 \\ c & 0 & x_1^2x_2^2 & x_1 & 0 & 0 & 0 & 0 & x_2 \\ d & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ e & 0 & 0 & 0 & x_1 & x_1^2 & 0 & 0 & x_2 \\ f & 0 & 0 & 0 & 0 & x_1^2 & x_1x_2^2 & x_2^2 & 0 \end{bmatrix}, \quad (12)$$

where we assume we are computing over \mathbb{Z}_2 .

As in standard homology, the boundary operator connects the graded chain modules into a chain complex C_* (Equation 3) and the i^{th} homology module is defined exactly as before (Equation 4):

$$H_i = \ker \partial_i / \text{im } \partial_{i+1} \quad (13)$$

4.2 Recasting the Problem

Our goal is to compute the homology modules. Following the definition, we have three tasks:

1. Compute the boundary module $\text{im } \partial_{i+1}$.
2. Compute the cycle module $\ker \partial_i$.
3. Compute the quotient H_i .

We next translate these three tasks into problems in computational commutative algebra. Both the boundary and cycle modules turn out to be submodules of free and finitely generated modules that consist of vectors of polynomials. For the rest of this paper, we assume that we are computing homology over the field k . Recall from Section 2.4 that our module is defined over the n -graded polynomial ring $A^n = k[x_1, \dots, x_n]$ with standard grading $A_v^n = kx^v, v \in \mathbb{N}^n$. For notational simplicity, we will use $R = A^n$ to denote this ring for the remainder of this section.

Let R^m be the Cartesian product of m copies of R . In other words, R^m consists of all column m -vectors of polynomials:

$$R^m = \{[f_1, \dots, f_m]^T \mid f_i \in R, 1 \leq i \leq m\}. \quad (14)$$

To distinguish elements of R^m from polynomials, we adopt the standard practice of placing them in bold format, so that $\mathbf{f} \in R^m$ is a vector of polynomials, but $f \in R$ is a polynomial. We use this practice exclusively for elements of R^m and not for other vectors, such as elements of \mathbb{N}^n . We now recast the three problems above into the language of computational commutative algebra.

1. The boundary module is a submodule of the polynomial module. The matrix for ∂_{i+1} has m_i rows and m_{i+1} columns, where m_j denotes the number of j -simplices in the complex. Let $F = (\mathbf{f}_1, \dots, \mathbf{f}_{m_{i+1}})$, $\mathbf{f}_i \in R^{m_i}$, where \mathbf{f}_i is the i th column in M_{i+1} . This tuple of polynomial vectors generate a submodule of R^{m_i}

$$\langle F \rangle = \left\{ \sum_{j=1}^{m_{i+1}} q_j \mathbf{f}_j \mid q_j \in R \right\}. \quad (15)$$

The *Submodule Membership Problem* asks whether a polynomial vector \mathbf{f} is in a submodule M ; That is, whether we may write \mathbf{f} in terms of some basis F as above. A solution to this problem would complete our first task.

2. The cycle submodule is also a submodule of the polynomial module. The matrix for ∂_i has m_{i-1} rows and m_i columns. Let $F = (\mathbf{f}_1, \dots, \mathbf{f}_{m_i})$, $\mathbf{f}_i \in R^{m_{i-1}}$, where \mathbf{f}_i is the i th column in the matrix. Given F , the set of all $[q_1, \dots, q_{m_i}]^T$, $q_i \in R$ such that

$$\sum_{i=1}^{m_i} q_i \mathbf{f}_i = \mathbf{0} \quad (16)$$

is a R -submodule of R^{m_i} called the (*first*) *syzygy module* of $(\mathbf{f}_1, \dots, \mathbf{f}_{m_i})$, denoted $\text{Syz}(\mathbf{f}_1, \dots, \mathbf{f}_{m_i})$. A set of generators for this submodule would complete our second task.

3. Our final task is simple, once we have completed the first two tasks. All we need to do is test whether the generators of the syzygy submodule, our cycles, are in the boundary submodule. As we shall see, the tools which allow us to complete the first two tasks also resolve this question.

4.3 Algorithms

In this section, we begin by reviewing concepts from commutative algebra that involve the polynomial module R^m . We then look at algorithms for solving the submodule membership problem and computing generators for the syzygy submodule. In our treatment, we follow Chapter 5 of Cox, Little, and O’Shea [6].

The *standard basis* for R^m is $\{\mathbf{e}_1, \dots, \mathbf{e}_m\}$, where \mathbf{e}_i is the *standard basis vector* with constant polynomials 0 in all positions except 1 in position i . We use the “top down” order on the standard basis vectors, so that $\mathbf{e}_i > \mathbf{e}_j$ whenever $i < j$. A *monomial* \mathbf{m} in R^m is an element of the form $x^u \mathbf{e}_i$ for some i and we say \mathbf{m} *contains* \mathbf{e}_i .

For algorithms, we need to order monomials in both R and R^m . For $u, v \in \mathbb{N}^n$, we say $u >_{\text{lex}} v$ if the vector difference $u - v \in \mathbb{Z}^n$, the leftmost nonzero entry is positive. The *lexicographic order* $>_{\text{lex}}$ is a total order on \mathbb{N}^n . For example, $(1, 3, 0) >_{\text{lex}} (1, 2, 1)$ since $(1, 3, 0) - (1, 2, 1) = (0, 1, -1)$ and the leftmost nonzero entry is 1. Now, suppose x^u and x^v are monomials in R . We say $x^u >_{\text{lex}} x^v$ if $u >_{\text{lex}} v$. This gives us a *monomial order* on R . We next extend $>_{\text{lex}}$ to a monomial order on R^m using the “position-over-term” (POT) rule: $x^u \mathbf{e}_i > x^v \mathbf{e}_j$ if $i < j$, or if $i = j$ and $x^u >_{\text{lex}} x^v$. Now, every element $\mathbf{f} \in R^m$ may be written, in a unique way, as a k -linear combination of monomials \mathbf{m}_i ,

$$\mathbf{f} = \sum_{i=1}^m c_i \mathbf{m}_i, \quad (17)$$

where $c_i \in k$, $c_i \neq 0$ and the monomials \mathbf{m}_i are ordered according to the monomial order. We now define:

- Each $c_i \mathbf{m}_i$ is a *term* of \mathbf{f} .
- The *leading coefficient* of \mathbf{f} is $\text{LC}(\mathbf{f}) = c_1 \in k$.
- The *leading monomial* of \mathbf{f} is $\text{LM}(\mathbf{f}) = \mathbf{m}_1$.
- The *leading term* of \mathbf{f} is $\text{LT}(\mathbf{f}) = c_1 \mathbf{m}_1$.

Example 5 Let $\mathbf{f} = [5x_1x_2^2, 2x_1 - 7x_3^3]^T \in R^2$. Then, we may write \mathbf{f} in terms of the standard basis (Equation 17):

$$\mathbf{f} = 5[x_1x_2^2, 0]^T - 7[0, x_3^3]^T + 2[0, x_1]^T \quad (18)$$

$$= 5x_1x_2^2\mathbf{e}_1 - 7x_3^3\mathbf{e}_2 + 2x_1\mathbf{e}_2. \quad (19)$$

From the second line, the monomials corresponding to sum (17) are $\mathbf{m}_1 = x_1x_2\mathbf{e}_1$, $\mathbf{m}_2 = x_3^3\mathbf{e}_2$, and $\mathbf{m}_3 = x_1\mathbf{e}_2$. The second term of \mathbf{f} is $7[0, x_3^3]$ and we have $\text{LC}(\mathbf{f}) = 5$, $\text{LM}(\mathbf{f}) = x_1x_2^2$, and $\text{LT}(\mathbf{f}) = 5x_1x_2^2$.

Finally, we extend division and least common multiple to monomials in R and R^m . Given monomials $x^u, x^v \in R$, if $x^u >_{\text{lex}} x^v$, then x^v divides x^u with quotient $x^u/x^v = x^{u-v}$. Moreover, we define $w \in \mathbb{N}^n$ by $w_i = \max(u_i, v_i)$ and define the monomial x^w to be the *least common multiple* of x^u and x^v , denoted $\text{LCM}(x^u, x^v) = x^w$. Next, given monomials $\mathbf{m} = x^u \mathbf{e}_i$ and $\mathbf{n} = x^v \mathbf{e}_j$ in R^m , we say \mathbf{n} divides \mathbf{m} iff $i = j$ and x^v divides x^u , and define the *quotient* to be $\mathbf{m}/\mathbf{n} = x^u/x^v = x^{u-v}$. In addition, we define

$$\text{LCM}(x^u \mathbf{e}_i, x^v \mathbf{e}_j) = \begin{cases} \text{LCM}(x^u, x^v) \mathbf{e}_i, & i = j, \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

Clearly, the LCM of two monomials is a monomial in R and R^m , respectively.

Example 6 Let $\mathbf{f} = [x_1, x_1 x_2]^T$ and $\mathbf{g} = [x_2, 0]^T$ be elements of R^2 . Then, the LCM of their leading monomials is:

$$\text{LCM}(\text{LM}(\mathbf{f}), \text{LM}(\mathbf{g})) = \text{LCM}(x_1 \mathbf{e}_1, x_2 \mathbf{e}_1) \quad (21)$$

$$= x_1 x_2 \mathbf{e}_1. \quad (22)$$

We now turn our attention to the submodule membership problem. Formally, given a polynomial vector \mathbf{f} and a set of t polynomial vectors F , we would like to know whether $\mathbf{f} \in \langle F \rangle$. We may divide \mathbf{f} by F using the division algorithm DIVIDE in Figure 2. After division, we have $\mathbf{f} = (\sum_{i=1}^t q_i \mathbf{f}_i) + \mathbf{r}$, so if the remainder $\mathbf{r} = \mathbf{0}$, then $\mathbf{f} \in \langle F \rangle$. This condition, however, is not necessary for modules over multivariate polynomials ($n > 1$) as we may get a non-zero remainder even when $\mathbf{f} \in \langle F \rangle$.

Let M be a submodule and $\text{LT}(M)$ be the set of leading terms of elements of M . A *Gröbner basis* is a basis $G \subseteq M$ such that $\langle \text{LT}(G) \rangle = \langle \text{LT}(M) \rangle$. For Gröbner bases, whenever $\mathbf{f} \in \langle F \rangle$, we always get $\mathbf{r} = \mathbf{0}$ after division of \mathbf{f} by $\langle F \rangle$, so we have solved the membership problem. The BUCHBERGER algorithm in Figure 3 computes a Gröbner basis G starting from any basis F . The *syzygy polynomial vector* or *S-polynomial* $S(\mathbf{f}, \mathbf{g}) \in R^m$ of \mathbf{f} and \mathbf{g} is

$$S(\mathbf{f}, \mathbf{g}) = \frac{\mathbf{h}}{\text{LT}(\mathbf{f})} \mathbf{f} - \frac{\mathbf{h}}{\text{LT}(\mathbf{g})} \mathbf{g}, \quad \text{where} \quad (23)$$

$$\mathbf{h} = \text{LCM}(\text{LM}(\mathbf{f}), \text{LM}(\mathbf{g})). \quad (24)$$

The BUCHBERGER algorithm utilizes S -polynomials to eliminate the leading terms of polynomial vectors and complete the given basis to a Gröbner basis.

A Gröbner basis generated by the algorithm is neither minimal nor unique. A *reduced Gröbner basis* is a Gröbner basis G such that for all $\mathbf{g} \in G$, $\text{LC}(\mathbf{g}) = 1$ and no monomial of \mathbf{g} lies in $\langle \text{LT}(G - \{\mathbf{g}\}) \rangle$. A reduced Gröbner basis is both minimal and unique. We may compute a reduced Gröbner basis by reducing each polynomial in G in turn, replacing $g \in G$ with the remainder of $\text{DIVIDE}(g, G - \{g\})$. Since the algorithm is rather simple, we do not present pseudo-code for it. The DIVIDE, BUCHBERGER, and the reduction algorithms together solve the submodule membership problem and, in turn, our first task of computing $\text{im } \partial_{i+1}$.

DIVIDE($\mathbf{f}, (\mathbf{f}_1, \dots, \mathbf{f}_t)$)

```

1   $\mathbf{p} \leftarrow \mathbf{f}, \mathbf{r} \leftarrow \mathbf{0}$ 
2  for  $i \leftarrow 1$  to  $t$ 
3      do  $q_i \leftarrow 0$ 
4  while  $\mathbf{p} \neq \mathbf{0}$ 
5      do if  $\text{LT}(\mathbf{f}_i)$  divides  $\text{LT}(\mathbf{p})$  for some  $i$ 
6          then  $q_i \leftarrow q_i + \text{LT}(\mathbf{p})/\text{LT}(\mathbf{f}_i)$ 
7               $\mathbf{p} \leftarrow \mathbf{p} - (\text{LT}(\mathbf{p})/\text{LT}(\mathbf{f}_i)) \mathbf{f}_i$ 
8          else  $\mathbf{r} \leftarrow \mathbf{r} + \text{LT}(\mathbf{p})$ 
9               $\mathbf{p} \leftarrow \mathbf{p} - \text{LT}(\mathbf{p})$ 
10 return  $((q_1, \dots, q_t), \mathbf{r})$ 

```

Figure 2: The DIVIDE algorithm divides $\mathbf{f} \in R^m$ by an t -tuple $(\mathbf{f}_1, \dots, \mathbf{f}_t)$, $\mathbf{f}_i \in R^m$ to get $\mathbf{f} = (\sum_{i=1}^m q_i \mathbf{f}_i) + \mathbf{r}$, where $q_i \in R$ and $\mathbf{r} \in R^m$.

```

BUCHBERGER( $F$ )
1  $G \leftarrow F$ 
2 repeat  $G' \leftarrow G$ 
3   foreach pair  $\mathbf{f} \neq \mathbf{g} \in G$ 
4      $((q_1, \dots, q_t), \mathbf{r}) \leftarrow \text{DIVIDE}(S(\mathbf{f}, \mathbf{g}), G)$ 
5     if  $\mathbf{r} \neq \mathbf{0}$ 
6       then  $G \leftarrow G \cup \{\mathbf{r}\}$ 
7   until  $G = G'$ 
8 return  $G$ 

```

Figure 3: The algorithm BUCHBERGER completes a given basis F to a Gröbner basis G by incrementally adding the remainders of S -polynomials, divided by the current basis.

We next compute generators for the syzygy submodule to complete our second task. We begin by computing a Gröbner basis $G = \{\mathbf{g}_1, \dots, \mathbf{g}_s\}$ for $\langle F \rangle$, where the vectors are ordered with some monomial order. We then compute $\text{DIVIDE}(S(\mathbf{g}_i, \mathbf{g}_j), G)$ for each pair of Gröbner basis elements. Since G is a Gröbner basis, the remainder of this division is $\mathbf{0}$, giving us

$$S(\mathbf{g}_i, \mathbf{g}_j) = \sum_{k=1}^s q_{ijk} \mathbf{g}_k. \quad (25)$$

Let $\epsilon_1, \dots, \epsilon_s$ be the standard basis vectors in R^s and let

$$\mathbf{h}_{ij} = \text{LCM}(\text{LT}(\mathbf{g}_i), \text{LT}(\mathbf{g}_j)), \quad (26)$$

$$\mathbf{q}_{ij} = \sum_{k=1}^s q_{ijk} \epsilon_k \in R^s. \quad (27)$$

For pairs (i, j) such that $\mathbf{h}_{ij} \neq 0$, we define $\mathbf{s}_{ij} \in R^s$ by

$$\mathbf{s}_{ij} = \frac{\mathbf{h}_{ij}}{\text{LT}(\mathbf{g}_i)} \epsilon_i - \frac{\mathbf{h}_{ij}}{\text{LT}(\mathbf{g}_j)} \epsilon_j - \mathbf{q}_{ij} \in R^s, \quad (28)$$

with $\mathbf{s}_{ij} = \mathbf{0}$, otherwise. *Schreyer's Theorem* states that the set $\{\mathbf{s}_{ij}\}_{i,j}$ form a Gröbner basis for $\text{Syz}(\mathbf{g}_1, \dots, \mathbf{g}_s)$ [6, Chapter 5, Theorem 3.3]. Clearly, we may compute this basis using DIVIDE. We use this basis to find generators for $\text{Syz}(\mathbf{f}_1, \dots, \mathbf{f}_t)$.

Let M_F and M_G be the $m \times t$ and $m \times s$ matrices in which the \mathbf{f}_i 's and \mathbf{g}_i 's are columns, respectively. As both bases generate the same module, there is a $t \times s$ matrix A and an $s \times t$ matrix B such that $M_G = M_F A$ and $M_F = M_G B$. To compute A , we simply initialize A to be the identity matrix and add a column to A for each division on line 4 of BUCHBERGER that records the pair involved in the S -polynomial. The matrix B may be computed by using the division algorithm. To see how, notice that each column of M_F is divisible by M_G since M_G is a Gröbner Basis for M_F ; now there is a column in B for each column $\mathbf{f}_i \in M_F$, which is obtained by division of \mathbf{f}_i by M_G . Let $\mathbf{s}_1, \dots, \mathbf{s}_t$ be the columns of the $t \times t$ matrix $I_t - AB$. Then,

$$\text{Syz}(\mathbf{f}_1, \dots, \mathbf{f}_t) = \langle A \mathbf{s}_{ij}, \mathbf{s}_1, \dots, \mathbf{s}_t \rangle, \quad (29)$$

giving us the syzygy generators [6, Chapter 5, Proposition 3.8]. We refer to the algorithm sketched above as *Schreyer's algorithm*. This algorithm completes our second task. The third task was to compute the quotient H_i given $\text{im } \partial_{i+1} = G$ and $\ker \partial_i = \text{Syz}(\mathbf{f}_1, \dots, \mathbf{f}_t)$. We simply need to find whether the columns of $\ker \partial_i$ can be represented as a combination of the columns of the $\text{im } \partial_{i+1}$. Now, H_i may be computed using the division algorithm. We divide every column in the $\ker \partial_i$ by the matrix $\text{im } \partial_{i+1}$ using the DIVIDE algorithm. If the remainder is non-zero, we add the remainder both to $\text{im } \partial_{i+1}$ and H_i (so as to count only unique cycles).

While the above algorithms solve the membership problem, they have not been used in practice due to their complexity. The submodule membership problem is a generalization of the *Polynomial Ideal Membership Problem (PIMP)* which is EXSPACE-complete, requiring exponential space and time [17, 20]. Indeed, the Buchberger algorithm, in its original form, is doubly-exponential and is therefore not practical.

5 Multigraded Algorithms

In this section, we show that multifiltrations provide additional structure that may be exploited to simplify the algorithms from the previous section. These simplifications convert these intractable algorithms into polynomial time algorithms.

5.1 Exploiting Homogeneity

The key property that we exploit for simplification is homogeneity.

Definition 6 (homogeneous) *Let M be an $m \times n$ matrix. The matrix M is homogeneous iff*

1. every column (row) \mathbf{f} of M is associated with a coordinate $u_{\mathbf{f}}$ and corresponding monomial $x^{u_{\mathbf{f}}}$,
2. every non-zero element M_{jk} may be expressed as the quotient of the monomials associated with column k and row j , respectively.

Any vector \mathbf{f} endowed with a coordinate $u_{\mathbf{f}}$ that may be written as above is homogeneous, e.g. the columns of M .

We will show that all boundary matrices ∂_i may be written as homogeneous matrices initially, and the algorithms for computing persistence only produce homogeneous matrices and vectors. That is, we maintain homogeneity as an invariant throughout the computation. We begin with our first task.

Lemma 1 *For a one-critical multifiltration, the matrix of $\partial_i: C_i \rightarrow C_{i-1}$ written in terms of the standard bases is homogeneous.*

Proof: Recall that we may write the boundary operator $\partial_i: C_i \rightarrow C_{i-1}$ explicitly as a $m_{i-1} \times m_i$ matrix M in terms of the standard bases for C_i and C_{i-1} , as shown in matrix (12) for ∂_1 . From Definition 5, the standard basis for C_i is the set of i -simplices in critical grades. In a one-critical multifiltration, each simplex σ has a unique critical coordinate u_{σ} (Definition (1)). In turn, we may represent this coordinate by the monomial $x^{u_{\sigma}}$. For instance, simplex a in Figure 1 has critical grade $(1, 1)$ and monomial $x^{(1,1)} = x_1x_2$. We order these monomials using $>_{\text{lex}}$ and use this ordering to rewrite the matrix for ∂_i . The matrix entry M_{jk} relates σ_k , the k th basis element for C_i to $\hat{\sigma}_j$, the j th basis element for C_{i-1} . If $\hat{\sigma}_j$ is not a face of σ_k , then $M_{jk} = 0$. Otherwise, $\hat{\sigma}_j$ is a face of σ_k . Since a face must precede a co-face in a multifiltration,

$$u_{\sigma_k} >_{\text{lex}} u_{\hat{\sigma}_j} \quad (30)$$

$$x^{u_{\sigma_k}} >_{\text{lex}} x^{u_{\hat{\sigma}_j}} \quad (31)$$

$$M_{jk} = x^{u_{\sigma_k}} / x^{u_{\hat{\sigma}_j}} = x^{u_{\sigma_k} - u_{\hat{\sigma}_j}}. \quad (32)$$

That is, the matrix is homogeneous. □

For example, $\hat{\sigma}_1 = a$ is a face of $\sigma_1 = ab$, so $M_{11} = x_1x_2^2/x_1x_2 = x_2$ in matrix (33).

Corollary 1 *For a one-critical multifiltration, the boundary matrix ∂_i in terms of the standard bases has monomial entries.*

Proof: The result is immediate from the proof of the previous lemma. The matrix entry is either 0, a monomial, or $x^{u(\sigma_k) - u(\hat{\sigma}_j)}$, a monomial. □

Example 7 We show the homogeneous matrix for ∂_1 below, where we augment the matrix with the associated monomials. Again, we assume we are computing over \mathbb{Z}_2 .

$$\left[\begin{array}{c|cccccccc} & ab & bc & cd & de & ef & af & bf & ce \\ & x_1x_2^2 & x_1^2x_2^2 & x_1 & x_1 & x_1^2 & x_1x_2^2 & x_2^2 & x_2 \\ \hline a & x_1x_2 & x_2 & 0 & 0 & 0 & x_2 & 0 & 0 \\ d & x_1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ b & 1 & x_1x_2^2 & x_1^2x_2^2 & 0 & 0 & 0 & x_2^2 & 0 \\ c & 1 & 0 & x_1^2x_2^2 & x_1 & 0 & 0 & 0 & x_2 \\ e & 1 & 0 & 0 & 0 & x_1 & x_1^2 & 0 & x_2 \\ f & 1 & 0 & 0 & 0 & 0 & x_1^2 & x_1x_2^2 & x_2^2 & 0 \end{array} \right] \quad (33)$$

We next focus on our second task, showing that given a homogeneous matrix as input, the algorithms produce homogeneous vectors and matrices. Let F be an $m \times n$ homogeneous matrix. Let $\{\mathbf{e}_1, \dots, \mathbf{e}_m\}$ and $\{\hat{\mathbf{e}}_1, \dots, \hat{\mathbf{e}}_n\}$ be the standard bases for R^m and R^n , respectively. A homogeneous matrix associates a coordinate and monomial to the row and column basis elements. For example, since x_1 is the monomial for row 2 of matrix (33), we have $u_{\mathbf{e}_2} = (1, 0)$ and $x^{u_{\mathbf{e}_2}} = x_1$. Each column \mathbf{f} in F is homogeneous and may be written in terms of rows:

$$\mathbf{f} = \sum_{i=1}^m c_i \frac{x^{u_{\mathbf{f}}}}{x^{u_{\mathbf{e}_i}}} \mathbf{e}_i, \quad (34)$$

where $c_i \in k$ and we allow $c_i = 0$ when a row is not used. For instance, column \mathbf{g} representing the edge ab in the bifiltration shown in Figure 1 may be written as:

$$\mathbf{g} = x_2 \mathbf{e}_1 + x_2 x_2^2 \mathbf{e}_3 \quad (35)$$

$$= \frac{x_2 x_2^2}{x_1 x_2} \mathbf{e}_1 + \frac{x_2 x_2^2}{1} \mathbf{e}_3 \quad (36)$$

$$= \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_1}}} \mathbf{e}_1 + \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_3}}} \mathbf{e}_3 = \sum_{i \in \{1,3\}} \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_i}}} \mathbf{e}_i, \quad (37)$$

Consider the BUCHBERGER algorithm in Figure 3. The algorithm repeatedly computes S -polynomials of homogeneous vectors on line 4.

Lemma 2 *The S -polynomial $S(\mathbf{f}, \mathbf{g})$ of homogeneous vectors \mathbf{f} and \mathbf{g} is homogeneous.*

Proof: A zero S -polynomial is trivially homogeneous. A non-zero S -polynomial $S(\mathbf{f}, \mathbf{g})$ implies that \mathbf{h} in Equation (24) is non-zero. By the definition of LCM (Equation 20), this implies that the leading monomials of \mathbf{f} and \mathbf{g} contain the same basis element \mathbf{e}_j . We have:

$$\text{LM}(\mathbf{f}) = \frac{x^{u_{\mathbf{f}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j \quad (38)$$

$$\text{LM}(\mathbf{g}) = \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j \quad (39)$$

$$\mathbf{h} = \text{LCM}(\text{LM}(\mathbf{f}), \text{LM}(\mathbf{g})) \quad (40)$$

$$= \text{LCM} \left(\frac{x^{u_{\mathbf{f}}}}{x^{u_{\mathbf{e}_j}}}, \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_j}}} \right) \mathbf{e}_j \quad (41)$$

$$= \frac{\text{LCM}(x^{u_{\mathbf{f}}}, x^{u_{\mathbf{g}}})}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j. \quad (42)$$

Let $x^\ell = \text{LCM}(x^{u_{\mathbf{f}}}, x^{u_{\mathbf{g}}}) = x^{\text{LCM}(u_{\mathbf{f}}, u_{\mathbf{g}})}$, giving us $\mathbf{h} = \frac{x^\ell}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j$. We now have

$$\frac{\mathbf{h}}{\text{LT}(\mathbf{f})} = \frac{\frac{x^\ell}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j}{c_{\mathbf{f}} \frac{x^{u_{\mathbf{f}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j} \quad (43)$$

$$= \frac{x^\ell}{c_{\mathbf{f}} x^{u_{\mathbf{f}}}}, \quad (44)$$

where $c_{\mathbf{f}} \neq 0$ is the field constant in the leading term of \mathbf{f} . Similarly, we get

$$\frac{\mathbf{h}}{\text{LT}(\mathbf{g})} = \frac{x^\ell}{c_{\mathbf{g}} x^{u_{\mathbf{g}}}}, \quad c_{\mathbf{g}} \neq 0 \quad (45)$$

Putting it together, we have

$$S(\mathbf{f}, \mathbf{g}) = \frac{\mathbf{h}}{\text{LT}(\mathbf{f})}\mathbf{f} - \frac{\mathbf{h}}{\text{LT}(\mathbf{g})}\mathbf{g} \quad (46)$$

$$= \frac{x^\ell}{c_{\mathbf{f}}x^{u_{\mathbf{f}}}} \sum_{i=1}^m c_i \frac{x^{u_{\mathbf{f}}}}{x^{u_{\mathbf{e}_i}}} \mathbf{e}_i - \frac{x^\ell}{c_{\mathbf{g}}x^{u_{\mathbf{g}}}} \sum_{i=1}^m c'_i \frac{x^{u_{\mathbf{g}}}}{x^{u_{\mathbf{e}_i}}} \mathbf{e}_i \quad (47)$$

$$= \sum_{i=1}^m d_i \frac{x^\ell}{x^{u_{\mathbf{e}_i}}} \mathbf{e}_i, \quad (48)$$

where $d_i = c_i/c_{\mathbf{f}} - c'_i/c_{\mathbf{g}}$. Comparing with Equation (34), we see that $S(\mathbf{f}, \mathbf{g})$ is homogeneous with $u_{S(\mathbf{f}, \mathbf{g})} = \ell$. \square

Having computed the S -polynomial, BUCHBERGER next divides it by the current homogeneous basis G on line 4 using a call to the DIVIDE algorithm in Figure 2.

Lemma 3 $\text{DIVIDE}(\mathbf{f}, (\mathbf{f}_1, \dots, \mathbf{f}_t))$ returns a homogeneous remainder vector \mathbf{r} for homogeneous vectors $\mathbf{f}, \mathbf{f}_i \in R^m$.

Proof: On line 1, \mathbf{r} and \mathbf{p} are initialized to be $\mathbf{0}$ and \mathbf{f} , respectively, and are both trivially homogeneous. We will show that each iteration of the **while** loop starting on line 4 maintains the homogeneity of these two vectors. On line 5, since both \mathbf{f}_i and \mathbf{p} are homogeneous, we have

$$\mathbf{f}_i = \sum_{j=1}^m c_{ij} \frac{x^{u_{\mathbf{f}_i}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j \quad (49)$$

$$\mathbf{p} = \sum_{j=1}^m d_j \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j. \quad (50)$$

Since $\text{LT}(\mathbf{f}_i)$ divides $\text{LT}(\mathbf{p})$, the terms must share basis element \mathbf{e}_k and we have

$$\text{LT}(\mathbf{f}_i) = c_{ik} \frac{x^{u_{\mathbf{f}_i}}}{x^{u_{\mathbf{e}_k}}} \mathbf{e}_k \quad (51)$$

$$\text{LT}(\mathbf{p}) = d_k \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{e}_k}}} \mathbf{e}_k \quad (52)$$

$$\text{LT}(\mathbf{p})/\text{LT}(\mathbf{f}_i) = \frac{d_k}{c_{ik}} \cdot \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{f}_i}}}, \quad (53)$$

where $x^{u_{\mathbf{p}}} >_{\text{lex}} x^{u_{\mathbf{f}_i}}$ so that the division makes sense. On line 7, \mathbf{p} is assigned to

$$\mathbf{p} - (\text{LT}(\mathbf{p})/\text{LT}(\mathbf{f}_i))\mathbf{f}_i = \sum_{j=1}^m d_j \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j - \left(\frac{d_k}{c_{ik}} \cdot \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{f}_i}}} \right) \sum_{j=1}^m c_{ij} \frac{x^{u_{\mathbf{f}_i}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j \quad (54)$$

$$= \sum_{j=1}^m \left(d_j - \frac{d_k \cdot c_{ij}}{c_{ik}} \right) \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j \quad (55)$$

$$= \sum_{j=1}^m d'_j \frac{x^{u_{\mathbf{p}}}}{x^{u_{\mathbf{e}_j}}} \mathbf{e}_j, \quad (56)$$

where $d'_j = d_j - d_k \cdot c_{ij}/c_{ik}$ and $d'_k = 0$, so the subtraction eliminates the k th term. The final sum means that \mathbf{p} is now a new homogeneous polynomial with the same coordinate $u_{\mathbf{p}}$ as before. Similarly, $\text{LT}(\mathbf{p})$ is added to \mathbf{r} on line 8 and subtracted from \mathbf{p} on line 9, and neither action changes the homogeneity of either vector. Both remain homogeneous with coordinate $u_{\mathbf{p}}$. \square

The lemmas combine to give us the desired result.

Theorem 1 (homogeneous Gröbner) *The BUCHBERGER algorithm computes a homogeneous Gröbner basis for a homogeneous matrix.*

Proof: Initially, the algorithm sets G to be the set of columns of the input matrix F , so the vectors in G are homogeneous by Lemma 1. On line 4, it computes the S -polynomial of homogeneous vectors $\mathbf{f}, \mathbf{g} \in G$. By Lemma 2, the S -polynomial is homogeneous. It then divides the S -polynomial by G . Since the input is homogeneous, DIVIDE produces a homogeneous remainder \mathbf{r} by Lemma 3. Since only homogeneous vectors are added to G on line 6, G remains homogeneous. We may extend this result easily to the reduced Gröbner basis. \square

Using similar arguments, we may show the following result, whose proof we omit here.

Theorem 2 (homogenous syzygy) *For a homogeneous matrix, all matrices encountered in the computation of the syzygy module are homogeneous.*

5.2 Data Structures and Optimizations

We have shown that the structure inherent in a multifiltration allows us to compute using homogeneous vectors and matrices whose entries are monomials only. We next explore the consequences of this restriction on both the data structures and complexity of the algorithms.

By Definition (6), an $m \times n$ homogeneous matrix naturally associates monomials to the standard bases for R^m and R^n . Moreover, every non-zero entry of the matrix is a quotient of these monomials as the matrix is homogeneous. Therefore, we do not need to store the matrix entries, but simply the field elements of the matrix along with the monomials for the bases. We may modify two standard data structures to represent the matrix.

- *linked list:* Each column stores its monomial as well as a linked-list of its non-zero entries in sorted order. The non-zero entries are represented by the row index and the field element. The matrix is simply a list of these columns in sorted order. Figure 4 displays matrix (33) in this data structure.
- *matrix:* Each column stores its monomial as well as the column of field coefficients. If we are computing over a finite field, we may pack bits for space efficiency.

The linked-list representation is appropriate for sparse matrices as it is space-efficient at the price of linear access time. This is essentially the representation used for computing in the one-dimensional setting [22]. In contrast, the matrix representation is appropriate for dense matrices as it provides constant access time at the cost of storing all zero entries. The multidimensional setting provides us with denser matrices, as we shall see, so the matrix representation becomes a viable structure.

In addition, the matrix representation is optimally suited to computing over the field \mathbb{Z}_2 , the field often commonly employed in topological data analysis. The matrix entries each take one bit and the column entries may be packed into machine words. Moreover, the only operation required by the algorithms is *symmetric difference* which may be implemented as a binary XOR operation provided by the chip. This approach gives us bit-level parallelism for free: On a 64-bit machine, we perform symmetric difference 64 times faster than on the list. The combination of these techniques allow the matrix structure to perform better than the linked-list representation in practice.

We may also exploit homogeneity to speed up the computation of new vectors and their insertion into the basis. We demonstrate this briefly using the BUCHBERGER algorithm. We order the columns of input matrix G using the POT rule for vectors as introduced in Section 4. Suppose we have $\mathbf{f}, \mathbf{g} \in G$ with $\mathbf{f} > \mathbf{g}$. If $S(\mathbf{f}, \mathbf{g}) \neq 0$, $\text{LT}(\mathbf{f})$ and




Figure 4: The linked list representation of the boundary matrix ∂_1 of Equation 12, for the bifiltration shown in Figure 1, in column sorted order. Note that the columns in Equation 12 are not ordered while they are sorted correctly here.

$\text{LT}(\mathbf{g})$ contain the same basis, which the S -polynomial eliminates. So, we have $S(\mathbf{f}, \mathbf{g}) < \mathbf{g} < \mathbf{f}$. This implies that when dividing $S(\mathbf{f}, \mathbf{g})$ by the vectors in G , we need only consider vectors that are smaller than \mathbf{g} . Since the vectors are in sorted order, we consider each in turn until we can no longer divide. By the POT rule, we may now insert the new remainder column here into the basis G . This gives us a constant time insertion operation for maintaining the ordering, as well as faster computation of the Gröbner basis.

5.3 Complexity

In this section, we give simple polynomial bounds on our multigraded algorithms. These bounds imply that we may compute multidimensional persistence in polynomial time.

Lemma 4 *Let F be an $m \times n$ homogeneous matrix of monomials. The Gröbner basis G contains $O(n^2m)$ vectors in the worst case. We may compute G using BUCHBERGER in $O(n^4m^3)$ worst-case time.*

Proof: In the worst case, F contains nm unique monomials. Each column $\mathbf{f} \in F$ may have any of the nm monomials as its monomial when included in the Gröbner basis G . Therefore, the total number of columns in G is $O(n^2m)$. In computing the Gröbner Basis, we compare all columns pairwise, so the total number of comparisons is $O(n^4m^2)$. Dividing the S -polynomial takes $O(m)$ time. Therefore, the worst-case running time is $O(n^4m^3)$. \square

In practice, the number of unique monomials in the matrix is lower than the worst case. In computing persistence, for example, we may control the number of unique monomials by ignoring close pairs of gradings. The following lemma bounds the basis size and running time in this case.

Lemma 5 *Let F be an $m \times n$ homogeneous matrix with h of unique monomials. The Gröbner basis G contains $O(hn)$ vectors and may be computed in time $O(n^3h^2)$.*

The proof is identical to the previous lemma.

Lemma 6 *Let F be an $m \times n$ homogeneous matrix of monomials and G be the Gröbner Basis of F . The syzygy module S for G may be computed using Schreyer's algorithm in $O(n^4m^2)$ worst-case time.*

Proof: In computing the syzygy Module, we compare all columns of G pairwise, so the total number of comparisons is $O(n^4m^2)$. Dividing the S -polynomial takes $O(m)$ time. Therefore, the worst-case running time is $O(n^4m^3)$. \square

Theorem 3 *Multidimensional persistence may be computed in polynomial time.*

Proof: Multidimensional persistence is represented by the Gröbner Bases and the syzygy moduli of all the homogeneous boundary matrices ∂_i for a given multifiltration. In the previous lemmas, we have shown that both the Gröbner Basis and the syzygy module can be computed in polynomial time. Therefore, one can compute multidimensional persistence in polynomial time. \square

In other words, our optimizations in this section turn the exponential-algorithms from the last section into polynomial-time algorithms.

6 Computing the Rank Invariant

Having described our algorithms, in this section we discuss the computation of the rank invariant. Recall that our solution is complete, but not an invariant. In contrast, the rank invariant (Equation 5) is incomplete, but is an invariant and may be used, for instance, as a descriptor in order to compare and match multifiltrations. We begin with a naive algorithm that computes the invariant for each pair independently. We then discuss alternate approaches using posets and vineyards. Direct computation, however, suffers from the storage requirement, which is intractable. Therefore, we end this section by giving an algorithm for reading off the rank invariant from the solution computed using our multigraded algorithms.

6.1 Direct Computation

We assume we are given a n -dimensional multifiltration of a cell complex K with m cells. Recall the rank invariant, Equation (5), from Section 2. We observe that any pair $u \leq v \in \mathbb{N}^n$ defines a one-dimensional filtration: We let t be a new parameter, map u to $t = 0$, v to $t = 1$, and obtain a two-level filtration. We then use the persistence algorithm to obtain barcodes [22]. The invariant $\rho_i(u, v)$ may be read off from the β_i -barcode: It is the number of intervals that contain both 0 and 1. The persistence algorithm is $\Theta(m^3)$ in the worst-case, so we have a cubic time algorithm for computing the rank invariant for a single pair of coordinates.

To fully compute the rank invariant, we need to consider all distinct pairs of complexes in a multifiltration. It may seem, at first, that we need to only consider critical coordinates, such as $(1, 1)$ and $(2, 0)$ in the bifiltration in Figure 1. However, note that the complex at coordinate $(2, 1)$ is also distinct even though no simplex is introduced at that coordinate. Inspired by this example, we may devise the following worst-case construction: We place m/n cells on each of the n axis to generate $(m/n)^n = \Theta(m^n)$ distinct complexes. Simple calculation shows that there are $\Theta(m^{2n})$ comparable coordinates with distinct complexes. For each pair, we may compute the rank invariant using our method above for a total of $O(m^{2n+3})$ running time. To store the rank invariant, we also require $\Theta(m^{2n})$ space.

6.2 Alternate Approaches

Our naive algorithm above computes the invariant for each pair of coordinates independently. In practice, we may read off multiple ranks from the same barcode for faster calculation. Any monotonically increasing path from the origin to the coordinate of the full complex is a one-dimensional filtration, such as the following path in Figure 1.

$$(0, 0) \rightarrow (1, 1) \rightarrow (2, 2) \rightarrow (3, 2) \quad (57)$$

Having computed persistence, we may read off the ranks for all six comparable pairs within this path. We may formalize this approach using language from the theory of *partially ordered sets*. The path described above is a *maximal chain* in the multifiltration poset: a maximal set of pairwise comparable complexes. We require a set of maximal chains such that each pair of comparable elements (here, complexes) are in at least one chain. Each maximal chain requires a single one-dimensional persistence computation. We now require an algorithm that computes the smallest set of such chains.

Another approach is to use Vineyards as introduced in [4]. The basic idea is the following: given two functions, compute a homotopy parameterized by λ ; for each λ compute the persistence barcodes and finally study the stack of these barcodes. A drawback of this approach is that a homotopy of functions might suffer from *range compression*. Consider a topological feature which exists for a long range in both the filtrations, but it is possible that the range for which the feature exists in the homotopy is rather small. Thus studying the vineyard, one would see this feature for a very short while, even though the feature existed for a long range in both the filtrations.

6.3 Multigraded Approach

Regardless of the approach, the direct computation is hampered by the exponential storage requirement, motivating our work in computing the rank invariant indirectly. Therefore, we first compute multidimensional persistence using our multigraded algorithms in Section 5. We then simply read off the rank invariant using the RANK algorithm, as shown in Figure 5. We describe the algorithm in the proof of the following theorem.

Theorem 4 $\text{RANK}(Z, B, u, v)$ computes the rank invariant $\rho_i(u, v)$, if Z is the syzygies of ∂_i and B is the Gröbner basis for ∂_{i+1} .

Proof: The algorithm uses two simple helper procedures. The procedure PROMOTE takes a matrix M and coordinate u as input. It then finds the columns $\mathbf{f} \in M$ whose associated coordinate $u_{\mathbf{f}}$ precedes u , and promotes them to coordinate u by a simple shift. The procedure QUOTIENT finds the quotient of the input matrices by division: If the remainder \mathbf{r} is non-zero, it adds \mathbf{r} to the quotient Q , also adding it to B so it only find unique cycles. Now assume the input are as in the statement of the theorem. By the definition of the rank invariant, we need to count homology cycles that exist at u and persist until v . The RANK algorithm implements this. We compute homology H_u at u on the first three lines. On line 4, we promote these cycles to coordinate v . We then quotient with boundaries B_v at v to find homology cycles H_{uv} that exist at u and persist until v . The size of this set is the rank invariant by definition. \square


```

RANK( $Z, B, u, v$ )
1  $Z_u \leftarrow \text{PROMOTE}(Z, u)$ 
2  $B_u \leftarrow \text{PROMOTE}(B, u)$ 
3  $H_u \leftarrow \text{QUOTIENT}(Z_u, B_u)$ 
4  $Z_{uv} \leftarrow \text{PROMOTE}(H_u, v)$ 
5  $B_v \leftarrow \text{PROMOTE}(B, v)$ 
6  $H_{uv} \leftarrow \text{QUOTIENT}(Z_{uv}, B_v)$ 
7 return  $|H_{uv}|$ 

QUOTIENT( $Z, B$ )
1  $Q \leftarrow \emptyset$ 
2 foreach  $\mathbf{f} \in Z$ 
3   do  $((q_1, \dots, q_t), \mathbf{r}) \leftarrow \text{DIVIDE}(\mathbf{f}, B)$ 
4     if  $\mathbf{r} \neq \mathbf{0}$ 
5       then  $Q \leftarrow Q \cup \{\mathbf{r}\}$ 
6          $B \leftarrow B \cup \{\mathbf{r}\}$ 
7 return  $Q$ 

PROMOTE( $M, u$ )
1 return  $\left\{ \frac{u}{u_{\mathbf{f}}} \mathbf{f} \mid \mathbf{f} \in M, u_{\mathbf{f}} \leq u \right\}$ 

```

Figure 5: The algorithm RANK computes the rank invariant $\rho_i(u, v)$ if Z is the set of syzygies of ∂_i and B is the Gröbner basis for ∂_{i+1} . The procedure QUOTIENT finds the quotient of Z by B using the DIVIDE algorithm. The procedure PROMOTE promotes cycles that exist before time u to that time.

7 Experiments

In this section, we describe our implementation as well as initial quantitative experiments that show the performance of our algorithms in practice. We end with a last look at our example bifiltration in Figure 1: computing its rank invariant using our multigraded algorithms.

7.1 Implementation

We initially used software packages CoCoA[3] and *Macaulay* [11], which contain standard implementations of the algorithms. These packages were immensely helpful during our software development as they allowed for quick and convenient testing of the basic algorithms. In practice, there are two problems in using these packages for large datasets. First, these packages are slow since they are general and not customized for homogeneous matrices. Second, these packages produce verbose output that must be parsed for further computation.

Our experience led us to implement our algorithms for computation over \mathbb{Z}_2 , optimizing the code for this field. Our implementation is in Java and was tested under Mac OS X 10.5.6 running on a 2.3Ghz Intel Quad-Core Xeon MacPro computer with 6GB RAM.

7.2 Data

We generate $n \times n$, random, bifiltered, homogeneous matrices, to simulate the boundary matrix ∂_{k-1} of a random bifiltered complex with n simplices in dimensions $k-1$ and $k-2$. We use the following procedure:

1. Randomly generate n monomials $\{m_1, \dots, m_n\}$ corresponding to the monomials associated with the basis elements of the rows.
2. For each column \mathbf{f} generate k integers indexing the non-zero rows.

3. Set the column monomial to be $\text{LCM}(m_j)$, where $\{m_j\}_j$ are the monomials of rows with non-zero

Each column in this matrix has k non-zero elements and is homogeneous by construction. We also generate random matrices but limit the number of unique monomials in the matrix to be $O(h^2)$ for different values of h . The basic idea behind these tests is that the range of the filtrations in a simplicial complex can typically be divided into smaller discrete intervals. For generation, we replace the first step of the procedure above with the following two steps:

0. Randomly generate h unique monomials $\{l_1, \dots, l_h\}$.

1. Generate n monomials $\{m_1, \dots, m_n\}$ corresponding to the monomials associated with the basis elements of the rows such that $m_i \in \{l_1, \dots, l_h\}$.

After executing step 2 and 3 above, our resulting matrix has homogeneous columns with k non-zero elements and at most h^2 unique monomials.

7.3 Size & Timings

According to Lemma 4, the number of columns in the Gröbner basis for a random matrix may grow $O(n^3)$ (we have $n = m$ here.) Figure 6(a) shows that the growth of the Gröbner basis is less in practice (about linear for $k = 2$ and quadratic for $k = 4$) and increases as the matrix becomes denser. Similarly, the theoretical running time for this matrix is $O(n^7)$. Figure 6(b) demonstrates that the actual running time matches this bound quite well (about $O(n^6)$ in these tests.) The matrix method, however, is considerably more efficient, about 20 times faster for our largest test here.




Figure 6: Random $n \times n$ matrices with k non-zero entries in each column. (a) The number of columns $|G|$ in Gröbner basis G (b) Running time in seconds for computing multidimensional persistence using list (l) or matrix (m) data structures.

We next limit the number of unique monomials in the input matrices. Figures 7 and 8 give the size and running time for matrices with at most h^2 monomials for $h = 20$ and $h = 100$, respectively. We see that the growth of the basis is about linear for different values of k and h , and the running time matches the theoretical $O(n^3)$ bound in Lemma 5 quite well.

7.4 Rank Invariant

We end this paper by revisiting our motivating bifiltration from Figure 1 and computing its multidimensional persistence and rank invariants using our algorithms. Using the natural ordering on the simplices, one can write the boundary

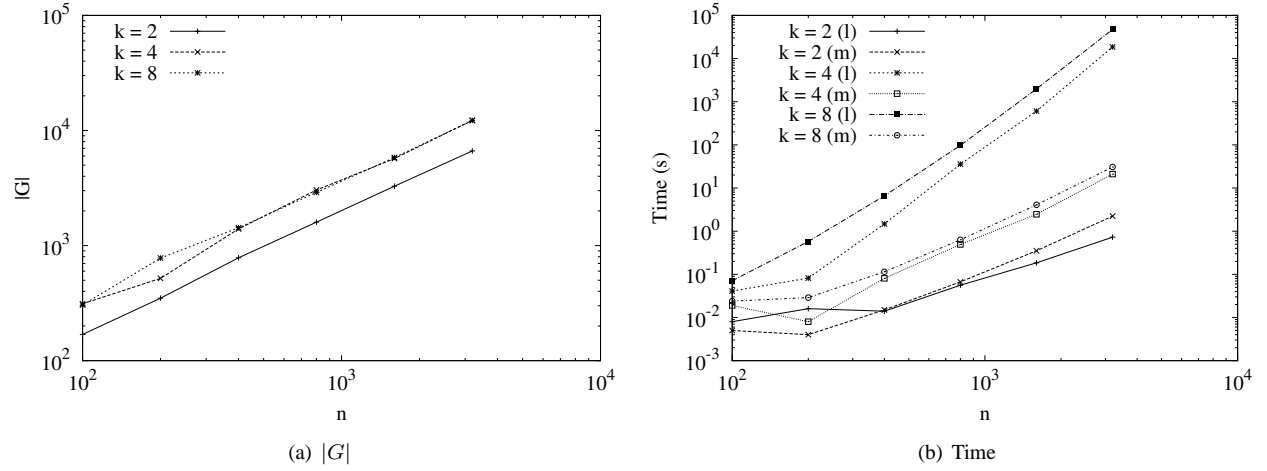


Figure 7: Random $n \times n$ matrices with k non-zero entries in each column and a total of h^2 monomials for $h = 20$. (a) The number of columns $|G|$ in Gröbner basis G . (b) Running time in seconds for computing multidimensional persistence using list (l) or matrix (m) data structures.

matrices as:

$$\partial_1 = \begin{bmatrix} 0 & x_1x_2^2 & 0 & 0 & 0 & x_2^2 & 0 & x_1^2x_2^2 \\ 0 & 0 & 0 & x_1 & 0 & 0 & x_2 & x_1^2x_2^2 \\ 0 & 0 & x_1 & 0 & x_1^2 & 0 & x_2 & 0 \\ x_1x_2^2 & 0 & 0 & 1 & x_1^2 & x_2^2 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ x_2 & x_2 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (58)$$

$$\partial_2 = \begin{bmatrix} 0 & x_1^3 \\ x_1^2 & 0 \\ 0 & x_1^2x_2 \\ 0 & x_1^2x_2 \\ x_1 & 0 \\ x_1 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \quad (59)$$

The Gröbner Basis (G_1) and the set of syzygies (Z_1) for ∂_1 are:

$$G_1 = \begin{bmatrix} 0 & x_1x_2^2 & 0 & 0 & 0 & x_2^2 & 0 & 0 & x_1^2x_2^2 \\ 0 & 0 & 0 & x_1 & 0 & 0 & x_1 & x_2 & x_1^2x_2^2 \\ 0 & 0 & x_1 & 0 & x_1^2 & 0 & x_1 & x_2 & 0 \\ x_1x_2^2 & 0 & 0 & 0 & x_1^2 & x_2^2 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ x_2 & x_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (60)$$

$$Z_1 = \begin{bmatrix} 0 & 0 & x_1 \\ x_1^2 & x_1 & 0 \\ x_1x_2^2 & 0 & x_2 \\ x_1x_2^2 & 0 & x_2 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ x_2^2 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \quad (61)$$

Note that each row in the syzygy matrix corresponds to an edge in the appropriate order. Finally, the Gröbner Basis for is $G_2 = \partial_2$ (since the only possible S -polynomial is identically 0).




Figure 8: Random $n \times n$ matrices with k non-zero entries in each column and a total of h^2 monomials for $h = 100$. (a) The number of columns $|G|$ in Gröbner basis G . (b) Running time in seconds for computing multidimensional persistence using list (l) or matrix (m) data structures.

Using G_1 , Z_1 and G_2 , one can read off the rank invariants for various u and v using the RANK algorithm in Section 6.3. A few interesting rank invariants for this example are:

u	v	$\rho_0(u, v)$	$\rho_1(u, v)$
[0, 0]	[1, 1]	3	0
[1, 0]	[2, 1]	2	0
[1, 1]	[1, 2]	2	1
[2, 2]	[3, 2]	1	1

8 Conclusion

In this paper, we give a full treatment to the computation of multidimensional persistence, from theory to implementation and experiments. We develop polynomial time algorithms by recasting the problem into computational commutative algebra. Although the recast problem is EXPSpace-complete, we exploit the multigraded setting to develop practical algorithms. The Gröbner bases we construct allow us to reconstruct the entire multidimensional persistence vector space, providing us a convenient way to compute the rank invariant. We implement all algorithms in the paper and show that the calculations are feasible due to the sparsity of the boundary matrices.

For additional speedup, we plan to parallelize the computation by batching and threading the XOR operations. We also plan to apply our algorithms toward studying scientific data. For instance, for zero-dimensional homology, multidimensional persistence corresponds to *clustering* multiparameterized data. This gives us a fresh perspective, as well as a new arsenal of computational tools, to attack an old and significant problem in data analysis.

References

- [1] CARLSSON, G., ISHKHANOV, T., DE SILVA, V., AND ZOMORODIAN, A. On the local behavior of spaces of natural images. *International Journal of Computer Vision* 76, 1 (2008), 1–12.
- [2] CARLSSON, G., AND ZOMORODIAN, A. The theory of multidimensional persistence. *Discrete & Computational Geometry* 42, 1 (2009), 71–93.
- [3] CoCoATEAM. CoCoA: a system for doing Computations in Commutative Algebra. <http://cocoa.dima.unige.it>.

- [4] COHEN-STEINER, D., EDELSBRUNNER, H., AND MOROZOV, D. Vines and vineyards by updating persistence in linear time. In *Proc. ACM Symposium on Computational Geometry* (2006), pp. 119–126.
- [5] COLLINS, A., ZOMORODIAN, A., CARLSSON, G., AND GUIBAS, L. A barcode shape descriptor for curve point cloud data. *Computers and Graphics* 28 (2004), 881–894.
- [6] COX, D. A., LITTLE, J., AND O’SHEA, D. *Using algebraic geometry*, second ed., vol. 185 of *Graduate Texts in Mathematics*. Springer, New York, 2005.
- [7] DE SILVA, V., AND CARLSSON, G. Topological estimation using witness complexes. In *Proc. Symposium on Point-Based Graphics* (2004), pp. 157–166.
- [8] DE SILVA, V., GHRIST, R., AND MUHAMMAD, A. Blind swarms for coverage in 2-D. In *Proceedings of Robotics: Science and Systems* (2005). <http://www.roboticsproceedings.org/rss01/>.
- [9] EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. Topological persistence and simplification. *Discrete & Computational Geometry* 28 (2002), 511–533.
- [10] FROSINI, P., AND MULAZZANI, M. Size homotopy groups for computation of natural size distances. *Bull. Belg. Math. Soc. Simon Stevin* 6, 3 (1999), 455–464.
- [11] GRAYSON, D. R., AND STILLMAN, M. E. Macaulay 2, a software system for research in algebraic geometry. <http://www.math.uiuc.edu/Macaulay2/>.
- [12] GROMOV, M. Hyperbolic groups. In *Essays in Group Theory*, S. Gersten, Ed. Springer Verlag, New York, NY, 1987, pp. 75–263.
- [13] GYULASSY, A., NATARAJAN, V., PASCUCCI, V., BREMER, P. T., AND HAMANN, B. Topology-based simplification for feature extraction from 3D scalar fields. In *Proc. IEEE Visualization* (2005), pp. 275–280.
- [14] HATCHER, A. *Algebraic Topology*. Cambridge University Press, New York, NY, 2002. <http://www.math.cornell.edu/~hatcher/AT/ATpage.html>.
- [15] LEVOY, M., AND HANRAHAN, P. Light field rendering. In *Proc. SIGGRAPH* (1996), pp. 31–42.
- [16] MATSUMOTO, Y. *An Introduction to Morse Theory*, vol. 208 of *Iwanami Series in Modern Mathematics*. American Mathematical Society, Providence, RI, 2002.
- [17] MAYR, E. W. Some complexity results for polynomial ideals. *Journal of Complexity* 13, 3 (1997), 303–325.
- [18] SINGH, G., MEMOLI, F., ISHKHANOV, T., SAPIRO, G., CARLSSON, G., AND RINGACH, D. L. Topological analysis of population activity in visual cortex. *Journal of Vision* 8, 8 (6 2008), 1–18.
- [19] TURK, G., AND LEVOY, M. Zippered polygon meshes from range images. In *Proc. SIGGRAPH* (1994), pp. 311–318.
- [20] VON ZUR GATHEN, J., AND GERHARD, J. *Modern Computer Algebra*, second ed. Cambridge University Press, Cambridge, UK, 2003.
- [21] ZHAO, F., AND GUIBAS, L. *Wireless Sensor Networks: An Information Processing Approach*. Morgan-Kaufmann, San Francisco, CA, 2004.
- [22] ZOMORODIAN, A., AND CARLSSON, G. Computing persistent homology. *Discrete & Computational Geometry* 33, 2 (2005), 249–274.