

Coreduction Homology Algorithm for Regular CW-Complexes

Paweł Dłotko

Institute of Computer Science
Jagiellonian University
ul. St. Łojasiewicza 6
30-348 Kraków, Poland

Tomasz Kaczynski

Département de mathématiques
Université de Sherbrooke
Sherbrooke, Québec, J1K 2R1, Canada

Marian Mrozek

Institute of Computer Science
Jagiellonian University
ul. St. Łojasiewicza 6
30-348 Kraków, Poland
and Division of Computational Mathematics
WSB-NLU
ul. Zielona 27
33-300 Nowy Sącz, Poland

Thomas Wanner

Department of Mathematical Sciences
George Mason University
Fairfax, VA 22030, USA

July 16, 2010

Abstract

In this paper we present a new algorithm for computing the homology of regular CW-complexes. This algorithm is based on the coreduction algorithm due to Mrozek and Batko and consists essentially of a geometric preprocessing algorithm for the standard chain complex generated by a CW-complex. By employing the concept of S-complexes the original chain complex can — in all known practical cases — be reduced to a significantly smaller S-complex with isomorphic homology, which can then be computed using standard methods. Furthermore, we demonstrate that in the context of non-uniform cubical grids this method significantly improves currently available algorithms based on uniform cubical grids.

AMS subject classifications: 55-04, 55N35, 52B99, 18G35, 57M20

Keywords: Homology algorithm, reduction methods, non-uniform cubical sets, S-complexes, CW-complexes

Contents

1	Introduction	3
2	S-Complexes and Elementary Coreductions	5
2.1	Preliminaries	5
2.2	S-Complexes and Regular S-Complexes	6
2.3	S-Reduction Pairs	8
2.4	Homology Generators	10
3	Coreduction Algorithm for CW-Complexes	10
3.1	CW-Complexes as S-Complexes	11
3.2	Reductions of Planar CW-Complexes	17
3.3	Rectangular CW-Complexes	20
4	Numerical Experiments	25
	References	31

1 Introduction

Homology has long been accepted as an important and computable tool for the analysis of topological spaces. The classical approach for computing homology is based on the celebrated Smith normal form, as described for example in [22]. For this method, computing the homology is basically reformulated as a linear algebra problem over the integers. The complexity of the Smith normal form algorithm, and consequently the homology algorithm based on Smith normal form, is $O(n^{3.376\dots})$, see [24]. This is sufficient in the case of spaces with small representation. Unfortunately, in modern scientific applications the size of the structures of interest frequently renders this approach infeasible due to the enormous size of the resulting matrices, and this fact is only exacerbated by modern data acquisition techniques. For example, in the context of materials science one might be interested in understanding the topology of extremely large complex microstructures, and computing their homology via the Smith normal form would easily lead to computations which cannot be handled on modern equipment. For some applications of topological methods in materials applications we refer the reader to [2, 10, 21, 25], as well as the references therein.

Based on these limitations of the Smith normal form algorithm, it is not surprising that more efficient algorithms for the computation of homology have been developed over the years, see for example [6, 7, 8, 9, 12, 13, 14, 19, 20], as well as the references therein. These algorithms and their implementations vary by the type of the requested input. Many algorithms require the topological space of interest represented as a simplicial complex, in some cases with certain dimensional restrictions. A different approach is employed in [12, 14, 19, 20], where the topological spaces are represented as cubical complexes, i.e., they are assembled from a finite collection of unit cubes. This framework is well-suited for the study of experimental and numerical data which is provided in pixel or voxel form, such as digital images or three-dimensional structures generated via experimental slicing techniques.

The success of many of these algorithms, for example [6, 19, 20], rests on minimizing or entirely avoiding the algebraization, i.e., the process of constructing the matrices of the boundary map followed by the Smith algorithm. One way of minimizing the algebraization consists in substantially reducing the representation of the topological space on input without changing its homology. A simple way to reduce the space is to collapse every face and its coboundary whenever the coboundary contains exactly one element. The collapsed space has the same homotopy type, so the homology does not change. Unfortunately, this method rarely leads to substantial reductions. However, the size of space significantly decreases after applying the dual concept of coreductions. The coreduction homology algorithm has been introduced in [19].

In the cubical setting, numerical experiments have shown that the coreduction algorithm is currently the fastest available method. In particular, it computes Betti numbers in a small fraction of the time used by other algorithms only to construct the




Figure 1: Two non-uniform cubical approximations: The left image contains a non-uniform grid that can be used to validate the nodal domain shown in dark blue, and the right image shows an efficient representation of a topological space with a singularity.

boundary matrices, i.e., even before they start applying the Smith algorithm. It has been proved in [11] that \mathbb{Z}_2 -Homology of weak 2-pseudomanifolds may be computed by a variant of the coreduction algorithm in $O(n\alpha(n))$, i.e., almost linear time. The coreduction algorithm is defined in the general framework of S-complexes, which will be described in more detail below.

Despite its efficiency, there are situations where the current form of the cubical coreduction algorithm does not allow for a timely homology computation. For example, in [4, 5] a rigorous computational technique was developed for finding the homology of nodal sets of smooth functions, a situation which is frequently encountered in applications. See for example [1, 16, 17] and the references therein. The method presented in [5], which will be explained in more detail in Section 4 below, is used to generate a non-uniform cubical grid which represents the considered nodal domain. An example of such a non-uniform grid is shown in the left image of Figure 1. Similarly, the right image in this figure shows an efficient representation of a topological space with a singularity.

While the method of [5] computes efficient non-uniform decompositions, the only way to compute the homology of the nodal sets using currently available homology codes is to embed the non-uniform grid into a fine uniform grid, whose size is determined by the smallest square in the non-uniform decomposition. As was pointed out in [5], this severely limits the practical applicability of the method. Similarly, in the context of computing the homology of maps it has been shown in [18] that the restriction to uniform cubical grids leads to inefficiencies which in principle can be addressed by representing the objects of interest in a less stringent way.

In the current paper we adapt the coreduction homology algorithm of [19] to the setting of regular CW-complexes. This allows us to prove that the Betti numbers of planar CW-complexes may be computed in linear time. We also introduce the concept

of a rectangular CW-complex of arbitrary dimension and show that the coreduction algorithm is particularly useful for such complexes.

Our method may be compared with the recent work [23] on computing the generators of image homology using graph pyramids. Although our main goal is somewhat different than the goal of [23], in the common setting of computing Betti numbers of 2-dimensional images the time and memory complexity of our algorithm is $O(p)$ with p the number of pixels, whereas the time complexity of building the graph pyramid alone is $O(p \log p)$. Therefore, in those applications where the graph pyramid is only a tool for finding Betti numbers and not part of a general goal, our approach to computing Betti numbers is superior. Our algorithm may also be adapted to compute the homology generators and in that case the time complexity is $O(pg)$ with g the number of generators. Since the complexity of the respective algorithm in [23] is at least $O(p \log p + g^2)$, the comparison in this case is in favor of our algorithm when the number of generators is bounded.

The remainder of the paper is organized as follows. In Section 2 we recall results and definitions from the theory of S-complexes and coreductions. Our main result on extending the coreduction method to regular CW-complexes is the subject of Section 3. We also demonstrate there that the necessary incidence indices can be determined easily in the non-uniform cubical setting of rectangular CW-complexes. Also here we show the complexity results concerning the planar case by proving that, in this case, the coreduction method always produces final S-complexes with trivial boundary operators, which eliminates the necessity to run a Smith normal form algorithm after the reduction. Finally, Section 4 contains some numerical examples.

2 S-Complexes and Elementary Coreductions

In this section we collect basic definitions and results for our underlying algebraic framework, i.e., the theory of S-complexes which was introduced in [19] as a reformulation of the concept of a free chain complex with a fixed basis in a manner suitable for computational reduction. Our discussion includes the notions of regular subsets of S-complexes, as well as of elementary coreductions and elementary reductions.

2.1 Preliminaries

We begin with fixing some terminology and notation. Throughout the paper we write \mathbb{Z} , \mathbb{R} , and R to denote the set of integers, the set of reals, and a fixed ring with unity, respectively. Given a set A we denote by $|A|$ its cardinality and by $R(A)$ the free module over R generated by A .

Given an object X of a category \mathcal{C} , a sequence $(X_q)_{q \in \mathbb{Z}}$ of objects of \mathcal{C} is a *gradation* of X , if X decomposes as the direct sum of the objects X_q .

An *abstract chain complex* is a free R -module C with a gradation $(C_n)_{n \in \mathbb{Z}}$ in the category of R -moduli and a collection of homomorphisms $\partial_n : C_n \rightarrow C_{n-1}$ called *boundary maps*, such that $\partial_{n-1}\partial_n = 0$ for all n . The elements of C_n are referred to as *n-chains* or simply *chains*. If $c \in C_n$ is a chain, then n is the *dimension* of c , denoted $\dim c$.

Given a basis S of C with gradation $(S_n)_{n \in \mathbb{Z}}$ we denote the associated scalar product of chains $c_1, c_2 \in C$ by $\langle c_1, c_2 \rangle_S$. We usually drop the index S if the basis S is clear from the context. The *incidence index* of basis elements $s, t \in S$ is defined as $\langle \partial s, t \rangle \in R$ if $\dim t = \dim s - 1$ and is zero otherwise. If the incidence index of s, t is nonzero, then t is called a *facet* of s . This defines a relation in $S \times S$, which will be referred to as *facet relation*. For any subset $A \subset S$ let us define

$$\begin{aligned} \text{bd}_S A &= \{ t \in S \mid t \text{ is a facet of } s \text{ for some } s \in A \}, \\ \text{cbd}_S A &= \{ s \in S \mid \text{there is a } t \in A \text{ such that } t \text{ is a facet of } s \}. \end{aligned}$$

The *weight* of a generator $s \in S$ is given by

$$w_S(s) := \max(|\text{bd}_S s|, |\text{cbd}_S s|),$$

and the *weight of the basis* S by

$$w(S) := \max \{ w_S(s) \mid s \in S \}.$$

Given a topological space X and a set $A \subset X$, we denote the closure and the boundary of A , respectively, by \bar{A} and ∂A .

2.2 S-Complexes and Regular S-Complexes

We first recall the definition of S-complexes as well as some basic results on the important special case of regular S-complexes.

Let S denote a finite set, and let $(S_q)_{q \in \mathbb{Z}}$ be a gradation of S with $S_q = \emptyset$ for all integers $q < 0$. Then the sequence $(R(S_q))_{q \in \mathbb{Z}}$ is a gradation of the module $R(S)$ in the category of moduli over the ring R . Let $\kappa : S \times S \rightarrow R$ be a map such that

$$\dim s = 1 + \dim t \quad \text{is satisfied whenever} \quad \kappa(s, t) \neq 0.$$

With κ we associate the map $\partial^\kappa : R(S) \rightarrow R(S)$ which is defined on generators $s \in S$ via

$$\partial^\kappa(s) = \sum_{t \in S} \kappa(s, t)t. \tag{1}$$

The pair (S, κ) is called an *S-complex* if $(R(S), \partial^\kappa)$ is a free chain complex with base S and *boundary map* ∂^κ .

For a given S-complex (S, κ) , its *homology* is defined as the homology of the chain complex $(R(S), \partial^\kappa)$, and it is denoted by $H(S, \kappa) = H(R(S), \partial^\kappa)$ or simply by $H(S)$. In the following, we will drop the superscript κ in ∂^κ whenever the meaning of κ is clear from the context.

One can easily observe that mathematically speaking an S-complex is just a chain complex with a fixed basis and the map κ provides the incidence indices with respect to this basis. The reason to introduce the terminology of S-complexes is to shift the emphasis from algebra to combinatorics. In the algebra setting, the moduli are given first and the bases are selected arbitrarily, leading to incidence indices obtained as respective entries in the matrices of boundary homomorphisms. In the setting of S-complexes we start with a collection of combinatorial objects (for instance simplices or cubes) for which a facet relation is somehow defined. Then, for each pair (s, t) of objects such that t is a facet of s we assign, in an algorithmic way, the incidence indices and use them to build the chain complexes. The shift of emphasis serves stressing some algorithmic aspects of our approach. As we will see in the sequel, in many situations only some or even no incidence indices are needed to compute the homology of an S-complex. Instead, the facet relation suffices. This fact may be used to speed up homology computations, because in many situations the facet relation is more straightforward and computationally cheaper than finding the incidence index. In particular, the computation of the incidence indices may be delayed so that only those indices are computed which are really needed.

However, to make this work we need a method of computing the incidence index after completing the reductions. We say that an S-complex is *effective*, if there is a formula or an algorithm which computes the incidence index $\kappa(s, t)$ by using the information in s and t only, i.e., without using the information in the whole S-complex. The two main examples of effective S-complexes are simplicial complexes and cubical complexes. In both of these cases one can readily derive explicit formulas for the incidence index $\kappa(s, t)$. For more details we refer the reader to [21]. As we will see in Section 3.1 a general CW-complex is not an effective S-complex. A detailed discussion of the case of non-uniform cubical complexes, which are effective S-complexes, will be given in Section 3.

In order to simplify a given S-complex through a reduction step, one is interested in replacing the original set of generators S by a subset $S' \subset S$, and the original incidence index κ by the restriction $\kappa' = \kappa|_{S' \times S'}$. This has to be done in such a way that (S', κ') is still an S-complex, that no evaluation of $\kappa(s, t)$ is needed to construct S' , and that $H(S) \cong H(S')$.

Characterizations of these requirements were obtained in [19]. We recall from there some terminology and the main result. A subset $S' \subset S$ is called *closed* in S if we have $\text{bd}_S S' \subset S'$, and it is called *open* in S if $S \setminus S'$ is closed in S . If $S' \subset S$ is such that for all $s, u \in S'$ and $t \in S$, $t \in \text{bd}_S s$ and $u \in \text{bd}_S t$ implies $t \in S'$, then S' is called *regular subset* of S . For further details consult [19]. A regular subset $T \subset S$ is called a *nullset*

of S , provided that T is closed or open in S and that $H(T) = 0$. The following result follows immediately from [19, Theorems 3.4 and 3.5].

Corollary 2.1 ([19], **Corollary 3.6**) *Let (S, κ) denote an S -complex over the ring R , and let $T \subset S$ be a nullset of S . Then the homologies $H(S)$ and $H(S \setminus T)$ are isomorphic.*

In other words, any nullset in an S -complex can be removed without changing the homology of S .

2.3 S-Reduction Pairs

We now describe a simple method for locating nullsets in an S -complex. This method is based on the concept of a reduction pair of a finitely generated free chain complex C , which was introduced in [13, 14] and which we now briefly recall. Given a free chain complex C with basis S we say that a pair (a, b) of elements of S is a *reduction pair* in C , if $\langle \partial b, a \rangle$ is invertible in the ring R . Notice that any reduction pair (a, b) satisfies $\dim b = 1 + \dim a$. For every such reduction pair one can define an associated chain complex $(\bar{C}, \bar{\partial})$ via

$$\bar{C}_q = \begin{cases} C_q & \text{for } q \notin \{\dim a, \dim b\}, \\ \{v \in C_q \mid \langle v, a \rangle = 0\} & \text{for } q = \dim a, \\ \{v \in C_q \mid \langle v, b \rangle = 0\} & \text{for } q = \dim b, \end{cases}$$

as well as

$$\bar{\partial}_q v = \begin{cases} \partial_q v & \text{for } q - 1 \notin \{\dim a, \dim b\}, \\ \partial_q v - \frac{\langle \partial_q v, a \rangle}{\langle \partial_q b, a \rangle} \partial_q b & \text{for } q - 1 = \dim a, \\ \partial_q v - \langle \partial_q v, b \rangle b & \text{for } q - 1 = \dim b. \end{cases}$$

In other words, one has to remove all chains from C which contain either a or b in their representation, and then update the boundary operator accordingly. One can show that the chain complexes (C, ∂) and $(\bar{C}, \bar{\partial})$ are chain equivalent. We call the chain complex $(\bar{C}, \bar{\partial})$ a *reduction* of (C, ∂) through the reduction pair (a, b) .

The fundamental philosophy behind reduction methods is the fact that in many cases it is advantageous to perform a sequence of reductions before applying the Smith normal form algorithm. In fact, experiments show that frequently the decrease in size of the chain complex is significant [19, 20]. From a computational point of view it is desirable to be able to identify reduction pairs quickly and then to be able to perform the reduction efficiently. This naturally leads to considering only specific types of reduction pairs. In this paper, we concentrate on the notions of *elementary reduction pairs* and *elementary coreduction pairs*, which were introduced in [19]. Let (S, κ) denote

an S-complex. Then a reduction pair (a, b) of elements of S is called an *elementary reduction pair* if we have $\text{cbd}_S a = \{b\}$, and in this case a is called a *free face* in S . Furthermore, the reduction pair (a, b) is called an *elementary coreduction pair* if we have $\text{bd}_S b = \{a\}$, and in this case b is called a *free coface* in S . Finally, we will use the term *S-reduction pair* to denote either an elementary reduction pair or an elementary coreduction pair. It turns out that S-reduction pairs give rise to nullsets in an S-complex, and we can therefore use them as the basis of a reduction algorithm for S-complexes which preserves homology. The following result is established in [19].

Theorem 2.2 ([19], **Theorem 4.1, Corollary 4.2**) *Let (S, κ) denote an S-complex over the ring R , and let $a, b \in S$. Then the following holds:*

- (a) *If (a, b) is an elementary reduction pair, then $\{a, b\}$ is open in S and a nullset.*
- (b) *If (a, b) is an elementary coreduction pair, then $\{a, b\}$ is closed in S and a nullset.*
- (c) *If (a, b) is an S-reduction pair, then the homologies $H(S)$ and $H(S \setminus \{a, b\})$ are isomorphic.*

Note that to detect an S-reduction pair only the knowledge of the facet relation is needed. In particular, there is no need to compute the incidence index.

The basic usage of reduction pairs can be described as follows. A *reduction sequence* of a chain complex C is a sequence of generator pairs $\alpha = \{(a_i, b_i)\}_{i=1,2,\dots,n}$ in S such that (a_i, b_i) is a reduction pair in C^{i-1} , where the chain complexes (C^i, ∂^i) are defined recursively by letting $(C^0, \partial^0) = (C, \partial)$, and then letting (C^i, ∂^i) denote the reduction of $(C^{i-1}, \partial^{i-1})$ through (a_i, b_i) , for $i = 1, 2, \dots, n$. We then use the notation $(C^\alpha, \partial^\alpha)$ for the last chain complex in the sequence $\{(C^i, \partial^i)\}_{i=1,2,\dots,n}$ and call this chain complex an α -*reduction* of (C, ∂) .

One can easily reformulate this concept in the framework of S-complexes. An *S-reduction sequence* of an S-complex (S, κ) is a sequence of S-reduction pairs $\alpha = \{(a_i, b_i)\}_{i=1,2,\dots,n}$ such that (a_i, b_i) is an S-reduction pair in (S^{i-1}, κ^{i-1}) , where the S-complexes (S^i, κ^i) are defined recursively by letting $(S^0, \kappa^0) = (S, \kappa)$, and then letting (S^i, κ^i) denote the reduction of (S^{i-1}, κ^{i-1}) through (a_i, b_i) , for $i = 1, 2, \dots, n$. We then use the notation $(S^\alpha, \kappa^\alpha)$ for the last S-complex in the sequence $\{(S^i, \kappa^i)\}_{i=1,2,\dots,n}$ and call this S-complex an α -*reduction* of (S, κ) .

Before closing this section, we would like to reiterate the main reason for introducing the concept of S-complexes and S-reduction pairs. At first glance it seems too restrictive to only consider these special cases of general reduction pairs. However, from the algorithmic point of view one has to be able to both quickly identify reduction pairs, and then to efficiently perform the reduction step. In the case of S-reduction pairs this does not require any algebraic operations. In particular, there is even no need to compute the incidence index. This contributes significantly to the speedup of homology computations.

2.4 Homology Generators

While Theorem 2.2 shows that performing a sequence of reductions via S-reduction pairs allows one to reduce the size of a chain complex without affecting its homology, even more is true. For every reduction step one can in fact explicitly write down the isomorphism guaranteed by Theorem 2.2. For this, let (a, b) denote a reduction pair. For the chain complex $(\bar{C}, \bar{\partial})$ the relevant chain maps are given by

$$\psi_k^{(a,b)}(c) = \begin{cases} c - \frac{\langle c, a \rangle}{\langle \partial b, a \rangle} \partial b & \text{for } k = \dim b - 1, \\ c - \langle c, b \rangle b & \text{for } k = \dim b, \\ c & \text{otherwise,} \end{cases}$$

and

$$\iota_k^{(a,b)}(c) = \begin{cases} c - \frac{\langle \partial c, a \rangle}{\langle \partial b, a \rangle} b & \text{for } k = \dim b, \\ c & \text{otherwise,} \end{cases}$$

and it is shown in [21, Theorem 2.8] that these chain maps are in fact mutually inverse chain equivalences. In the case of an S-reduction pair (a, b) these maps can be used also for the chain complex (C', ∂') obtained from the regular subset $S' = S \setminus \{a, b\}$ via the definitions $C' = R(S')$ and $\partial' = \partial^\kappa|_{C'}$. In fact, it is shown in [21, Theorem 2.9] that the chain complexes $(\bar{C}, \bar{\partial})$ and (C', ∂') coincide.

Given a reduction sequence α in S we let

$$\iota^\alpha = \iota^{(a_1, b_1)} \circ \iota^{(a_2, b_2)} \circ \dots \circ \iota^{(a_n, b_n)}, \quad (2)$$

$$\psi^\alpha = \psi^{(a_n, b_n)} \circ \psi^{(a_{n-1}, b_{n-1})} \circ \dots \circ \psi^{(a_1, b_1)}. \quad (3)$$

The chain maps ι^α and ψ^α described above allow one easily to track generators as well. Namely, once the representatives of homology generators are computed in the complex S^α , one can use the map ι^α to find the representatives of the homology generator in the initial complex S^0 . Therefore storing the elementary reduction and coreduction pairs enables to obtain information about the representatives of homology generators. Moreover, in the case of S-complexes, the cost of finding the image of a generator in the map ι^α is $O(|S|w(S))$, see [21, Theorem 3.1].

3 Coreduction Algorithm for CW-Complexes

In this section we describe how S-reduction pairs can be employed in the context of CW-complexes, and demonstrate its applicability specifically for the case of planar CW-complexes, as well as for nonuniform cubical complexes in arbitrary dimensions. In particular, for planar CW-complexes we show that the coreduction algorithm completely reduces the underlying S-complex, thereby eliminating the need for the Smith

normal form algorithm altogether. For nonuniform rectangular complexes we show that the incidence indices of two cubes can be determined immediately from the representation of the involved rectangular boxes.

3.1 CW-Complexes as S-Complexes

We begin by recalling basic results on CW-complexes, including a discussion of regular CW-complexes and incidence indices. For this, we essentially follow the presentation in [15]. First, given a Hausdorff space X , an n -cell e is a subset of X homeomorphic to the open n -dimensional unit ball $B_1^n \subset \mathbb{R}^n$. We call n the *dimension* of e and write $\dim e = n$. An n -cell e is *attached to a closed subset* $K \subset X$, if $e \cap K = \emptyset$ and there is a continuous map $f : \overline{B_1^n} \rightarrow \bar{e}$, called *characteristic map*, sending the open ball B_1^n homeomorphically onto e and such that $f(\partial B_1^n) \subset K$.

For the purposes of this paper, we only consider finite CW-complexes which are defined as follows. A closed subset $K \subset X$ is a *finite CW-complex* of dimension N , if there exists an ascending sequence of closed subspaces $K^0 \subset K^1 \subset \dots \subset K^N = K$ such that K^0 is a finite set consisting of 0-cells, also called *vertices*, and K^n is obtained from K^{n-1} by attaching a finite collection of n -cells for $n = 1, \dots, N$. The collection of n -cells of K is denoted by \mathcal{K}_n . We put $\mathcal{K}_n := \emptyset$ for $n \notin \{0, 1, \dots, N\}$. The graded set $\mathcal{K} := (\mathcal{K}_n)_{n \in \mathbb{Z}}$ is called the *CW-complex structure* of K . The subset K^n of X is the *n -skeleton*. An N -dimensional CW-complex is *regular*, if for each cell $e \in \mathcal{K}_n$, where $n = 0, 1, \dots, N$, there exists a characteristic map $f : \overline{B_1^n(0)} \rightarrow \bar{e}$ which is a homeomorphism. In this case, given two cells e, e' we say that e' is a *face* of e and write $e' < e$, if $e' \neq e$ and the inclusion $e' \subset \bar{e}$ holds. If $e' < e$ and $\dim e' = \dim e - 1$, then we say that e' is a *facet* of e .

One of the main reasons of the importance of CW-complexes is that they represent a wider class of spaces than simplicial and cubical sets, while preserving enough combinatorial structure to define the associated chain complex and homology in the combinatorial spirit of simplicial and cubical complexes. For the detailed definition of the chain complex structure associated with a CW-complex K we refer the reader to [15]. In the sequel we refer to this chain complex as the *standard chain complex* of K and denote it by $C(K)$. Here we only recall a few central ideas which will be used later. The first fact is that the homology of $C(K)$ coincides with the singular homology of K , so that $C(K)$ may be used to compute the homology of K . Next, the group of n -chains of $C(K)$ is generated by \mathcal{K}_n , so that we can consider the CW-complex structure $\mathcal{K} := (\mathcal{K}_n)_{n \in \mathbb{Z}}$ as an S-complex. However, to make this approach computationally useful, we need a method of determining the incidence indices $\kappa(e, f)$. For general CW-complexes the standard way of obtaining these is either by computing the Brouwer degree of an appropriate map as described in [3] or, equivalently, by computing relative homology. Unfortunately, neither of these methods is feasible in the context of a computational approach.

However, if one restricts attention to the special case of regular CW-complexes, the situation is different. Recall from [15] that for every pair $(f, e) \in \mathcal{K}_n \times \mathcal{K}_{n+2}$ of a regular CW-complex \mathcal{K} such that $f < e$, there are exactly two $(n+1)$ -cells g_k , with $k = 1, 2$, such that $f < g_k < e$. In the case of regular CW-complexes the following characterization of incidence indices, following immediately from [15, Theorem IX.7.2] is useful in the computational context.

Theorem 3.1 *Let K be a finite regular CW-complex on the Hausdorff space X . For each pair $(e, f) \in \mathcal{K}_n \times \mathcal{K}_{n-1}$ let there be given an integer $\alpha_{ef} \in \{0, \pm 1\}$ such that the following four conditions hold:*

- (a) *If f is not a facet of e , then $\alpha_{ef} = 0$.*
- (b) *If f is a facet of e , then $\alpha_{ef} = \pm 1$.*
- (c) *If f and g are the two vertices of the 1-cell e , then $\alpha_{ef} + \alpha_{eg} = 0$.*
- (d) *Let $e \in \mathcal{K}_n$ and $f \in \mathcal{K}_{n-2}$ be two cells of K such that $f < e$. Furthermore, let $g_1, g_2 \in \mathcal{K}_{n-1}$ denote the unique $(n-1)$ -cells such that $f < g_i < e$ for $i = 1, 2$. Then the identity $\alpha_{eg_1}\alpha_{g_1f} + \alpha_{eg_2}\alpha_{g_2f} = 0$ holds.*

Then, defining the boundary operator by (1) with

$$\kappa(e, f) := \begin{cases} \alpha_{ef} & \text{if } f \text{ is a face of } e, \\ 0 & \text{otherwise.} \end{cases}$$

leads to a well-defined chain complex isomorphic to the standard chain complex of K .

Theorem 3.1 enables the computation of incidence indices for any regular CW-complex. In fact, Theorem 3.1 leads directly to an algorithm for the computation of incidence indices of arbitrary regular CW-complexes. The algorithm is presented in Table 1.

The algorithm provides a means for the direct Smith normal form computation for arbitrary regular CW-complexes. Moreover, it follows that reduction and coreduction methods can be applied to speed up the homology computations. We would like to point out that the incidence indices provided by the algorithm in Table 1 can be obtained by one of the standard methods described for instance in [3]. The correctness of the algorithm in Table 1 is established in the following theorem.

Theorem 3.2 *For an arbitrary regular CW-complex structure \mathcal{K} the algorithm in Table 1 returns a collection of indices α_{ef} for f a facet of e , such that the assumptions of Theorem 3.1 are satisfied.*

- | |
|---|
| <ol style="list-style-type: none"> 1. for every 1-dimensional cell e in the complex: <ol style="list-style-type: none"> (a) Let f and g be the facets of e. Set $\alpha_{ef} := 1$, $\alpha_{eg} := -1$; 2. for $i = 2$ to the maximal dimension of cells in the complex: <ol style="list-style-type: none"> (a) for every i-cell e in the complex <ol style="list-style-type: none"> i. pick any f which is a facet of e and set $\alpha_{ef} := 1$; ii. Let Q be an empty queue; iii. push(Q, f); iv. while $Q \neq \emptyset$: <ul style="list-style-type: none"> • $f := \mathbf{pop}(Q)$; • for every facet g of e such that α_{eg} is not yet set and f and g are sharing a facet $h \in \mathcal{K}_{i-2}$: <ul style="list-style-type: none"> – Set $\alpha_{eg} := -\alpha_{ef}\alpha_{fh}/\alpha_{gh}$ – push(Q, g); |
|---|

Table 1: An algorithm for determining the incidence indices of cells in regular CW-complex which is based on Theorem 3.1.

Proof: By [15, Lemma IX.7.1], nonzero incidence index can only occur for a pair of a cell and its facet. Therefore, the algorithm only considers pairs of this type. If the value α_{ef} was not set by the algorithm, it is assumed to be 0. Moreover, it is straightforward to see that for the incidence indices between one- and zero-dimensional cells as defined in line (1.a) of the algorithm the condition (c) of Theorem 3.1 holds.

We now turn our attention to i -dimensional cells with $i \geq 2$, and assume that all lower-dimensional cells have already been taken care of. Let $i \geq 2$ be fixed and consider an i -cell e . Moreover, let $\mathcal{E}(e)$ denote the set of all equations in condition (d) in Theorem 3.1 which contain e as the highest-dimensional cell. In order to establish the theorem, one only has to show that for every facet g of the given cell e the variable α_{eg} is set by the algorithm in such a way that the equations in $\mathcal{E}(e)$ hold.

In line (2.a.i) of the algorithm, a facet f of e is (arbitrarily) distinguished. Setting the incidence index α_{ef} equal to 1 is easily seen to be equivalent to choosing the orientation of the given i -cell e . We would like to point out that according to [15, Section IX.5], this choice of orientation uniquely determines all incidence indices of the form α_{eg} , where g denotes a facet of e . In particular, this shows that the system $\mathcal{E}(e)$ of equations in condition (d) in Theorem 3.1 which involve e and lower-dimensional cells has a unique solution satisfying $\alpha_{ef} = 1$. See also [15, Lemma IX.7.1].

Now let us consider lines (2.ii-iv). As we mentioned in the previous paragraph, it

follows from [15, Section IX.5, Lemma IX.7.1] that the solution of the system $\mathcal{E}(e)$ exists and is unique, due to the choice made in line (2.i). In course of executing lines (2.ii-iv), the algorithm solves a subset $\mathcal{E}_{alg}(e) \subset \mathcal{E}(e)$ of these equations. Moreover, it follows from the structure of the algorithm that as the loop in (2.iv) is executed, every equation in $\mathcal{E}_{alg}(e)$ has a unique solution. If we now assume for the moment that the execution of lines (2.ii-iv) sets the values α_{eg} for all facets g of e , then these values in fact have to solve all equations in $\mathcal{E}(e)$ — the algorithm automatically avoids unnecessary equations in this overdetermined, yet consistent, system.

To complete the proof of the theorem we still have to show that the execution of lines (2.i-iv) in the algorithm actually determines the incidence indices α_{eg} for all facets g of e . For this, let \mathcal{P} denote the set of all facets g of e for which the incidence indices α_{eg} has been set by the algorithm, and let \mathcal{N} denote the set of the remaining facets of e . Suppose that in fact we have $\mathcal{N} \neq \emptyset$. The definition of the algorithm then implies that for every cell $g \in \mathcal{P}$, every cell $g' \in \mathcal{N}$, and every $(i - 2)$ -dimensional face f of e , the cell f cannot be a common facet of both g and g' . Consequently, the equations in $\mathcal{E}(e)$ from condition (d) in Theorem 3.1 can be partitioned into two sets, one which only provides relations involving cells in \mathcal{P} , and another one which only concerns cells in \mathcal{N} . This, however, implies that the initial cell $g' \in \mathcal{N}$ can be (arbitrarily) picked and the incidence index $\alpha_{eg'}$ can be set either to 1 or -1 . Then the **while** loop in line (2.iv) of the algorithm can be rerun. Thus, by possibly repeating the above procedure several times, we can construct at least two distinct solutions to the set of equations $\mathcal{E}(e)$ which satisfy $\alpha_{eg} = 1$ (one of them with $\alpha_{eg'}$ set to 1, another with $\alpha_{eg'}$ set to -1) — and this of course contradicts the uniqueness of the solution mentioned above. In other words, we have to have $\mathcal{N} = \emptyset$, and this completes the proof of the theorem. \diamond

Theorem 3.3 *If the implementation of the algorithm in Table 1 uses pointers to access the elements in the boundary or coboundary of a cell, then the time complexity of the algorithm is $O(|\mathcal{K}|w(\mathcal{K})^4)$.*

Proof: First observe that in the **for**-loop in line (1) and in the nested **for**-loops in lines (2), and (2.a) of the algorithm each cell from the complex \mathcal{K} is considered just once. The number of iterations of the **for**-loops in line (1) and in lines (2), and (2.a) is therefore $O(|\mathcal{K}|)$. The actions taken in the points *i.*, *ii.* and *iii.* of the algorithm, due to the used pointer data structure, require constant time. There remains to calculate the computational cost of the **while**-loop in the point *iv.* of the algorithm. There are at most $w(\mathcal{K})$ elements g being the facets of e . Each such element g will be considered in the **while**-loop in the point *iv.* of the algorithm just once. The total number of iterations of **while** loop is therefore bounded by $w(\mathcal{K})$. It is straightforward to observe that the total number of iterations of the **for**-loop inside the **while**-loop is $O(w(\mathcal{K})^3)$, therefore the total complexity of the **while**-loop is $O(w(\mathcal{K})^4)$. Consequently, the time complexity of the whole algorithm is $O(|\mathcal{K}|w(\mathcal{K})^4)$, as required. \diamond




Figure 2: Example 3.4 demonstrates the effect of S-coreductions on a simple regular CW-complex. Notice that in order to start the coreduction sequence we add the empty set in dimension -1 . (White circles indicate missing vertices.)

Note that in most applications the number $w(\mathcal{K})$ is constant and small when compared with $|\mathcal{K}|$ and consequently, in such situations, the total complexity of the algorithm is linear with respect to the number of cells in the CW-complex.

Example 3.4 In order to illustrate the S-reduction procedure in the context of regular CW-complexes, consider the simple CW-complex shown in the upper left part of Figure 2. Upon adding the empty set in dimension -1 of the associated chain complex (as indicated), one can apply a coreduction to remove one of the vertices of this complex. This yields the S-complex in the upper right part. Subsequently, a sequence of edge-vertex coreductions removes all the remaining vertices, as well as all the edges in a spanning tree of the 1-skeleton, resulting in the S-complex on the lower left. Finally, after two more coreductions between a 2-cell and an edge, one obtains the final complex shown in the lower right part of Figure 2. Notice that this S-complex has empty chain groups in dimensions 0 and 2, so the homology can be read off without any further computations: It is the homology of a pointed circle, i.e., it equals R in dimension 1 and is 0 otherwise. We will see in the next section that this is not a coincidence.

Before closing this section, let us point out once more that in order for the above coreduction algorithm to provide a significant improvement we need to be able to determine the incidence index κ in an efficient and easy way, unless the final chain complex is trivial. In the next two sections, these issues will be discussed in more detail. Notice, however, that it is in general not possible to determine κ by only satisfying the equations of Theorem 3.1 for the cells in the final S-complex.

Example 3.5 Consider for example the standard torus, represented by the regular CW-complex shown on the left of Figure 3. As usual, vertices and edges marked with




Figure 3: Example 3.5 demonstrates the necessity of a-priori knowledge of the incidence index κ . For the torus representation shown on the left, the coreduction algorithm results in the S-complex shown in the middle, which contains no vertices any more. Thus, the equations of Theorem 3.1 do not pose any restrictions, and the boundary operator defined in the picture gives the wrong homology.

the same letters are being identified.

By adding the empty set as before and performing coreductions as far as possible, one can easily see that the final S-complex is basically given by the complex shown in the middle of Figure 3. This complex consists of the five two-cells A , B , C , D , and E , as well as the six one-cells α , β , γ , δ , ϵ , and ζ . Yet, the coreduction process removed all zero-cells from the S-complex. Thus, in order to find the incidence indices which satisfy all the conditions (a)-(d) of Theorem 3.1, one only has to make sure that (a) and (b) are satisfied; the equations in parts (c) and (d) cannot occur. In particular, if we choose the incidence index κ via

κ	A	B	C	D	E
α	-1	0	0	0	+1
β	+1	0	-1	0	0
γ	0	0	+1	0	-1
δ	0	+1	0	+1	0
ϵ	0	-1	+1	0	0
ζ	0	0	+1	-1	0

then all restrictions in conditions (a)-(d) in Theorem 3.1 are satisfied. This choice of incidence index results in the boundary operator ∂ presented in the right part of Figure 3. However, one can easily check that over the ring of integers, the S-complex shown in the middle of the figure, equipped with this boundary operator ∂ , has the trivial homology groups $H_0 = 0$ and $H_2 = 0$, as well as the nontrivial one $H_1 = \mathbb{Z} \oplus \mathbb{Z}_2$. Due to the presence of torsion, this does clearly not furnish the (reduced) homology of the torus.




Figure 4: This example shows that in general it is not possible to completely reduce a planar CW-complex using only elementary reductions.



Figure 5: This example shows that in general it is not possible to completely reduce a planar CW-complex using only elementary coreductions. In order to start the coreduction sequence from the CW-complex shown on the left we add the empty set in dimension -1 .

The last example shows that a general CW-complex cannot be considered as an effective S-complex. This, of course, does not mean that the S-reductions cannot be applied to such a general complex. However, in such a case all incidence indices have to be determined by applying the algorithm in Table 1 before applying any S-reduction. But, as we will see in Section 3.3, there is an important case of CW-complexes, which is effective.

3.2 Reductions of Planar CW-Complexes

In this section we describe a situation in which S-reductions allow the complete reduction of the original S-complex. In this situation, one can determine the homology of the underlying CW-complex by just counting the number of remaining generators. It will turn out that this can always be achieved for planar regular CW-complexes.

Before we turn to the proof of this result, notice that in order to achieve such a complete reduction one usually has to employ both elementary reductions and elementary coreductions. To see this, Figure 4 shows a planar CW-complex and a reduction sequence which uses only elementary reductions. For the final complex shown on the right, no further elementary reductions are possible. Similarly, Figure 5 shows a planar CW-complex and a reduction sequence which uses only elementary coreductions. Also here, the final S-complex does not allow for any further elementary coreductions.

In order to completely reduce the S-complex associated with a planar regular CW-complex a few preparations are necessary. It was already mentioned at the end of the last section that in order to start the coreduction sequence, one has to add the empty set to the associated chain complex in dimension -1 . Thus, an elementary coreduction

can be used to remove one vertex from the S-complex. In a homological sense this of course means that one computes the reduced homology groups, rather than the original one. It is therefore not too hard to see that for a disconnected CW-complex, one has to add several elements to the associated chain complex in dimension -1 , namely one for each component, in order to finally arrive at a trivial zero-dimensional homology group. From an algorithmic point of view, this can easily be achieved. One simply removes a vertex¹, and applies elementary coreductions as long as possible. If the resulting S-complex still has generators in dimension 0, one just repeats the process, and so on. By keeping track how many times this process has to be started, one can keep track of the number of components in the underlying CW-complex.

We turn our attention to the main result of this section. For this, we call a regular CW-complex *planar*, if it can be embedded into \mathbb{R}^2 . The following result now shows that every such complex can be completely reduced by using only elementary S-reduction pairs.

Theorem 3.6 *Let K denote a finite planar regular CW-complex and let \mathcal{K} denote the associated CW-complex structure. Based on the discussion preceding this theorem, we may assume without loss of generality that K is connected. Then there exists an S-reduction sequence α of \mathcal{K} such that the α -reduction of \mathcal{K} has trivial chain groups except in dimension one. Furthermore, the cardinality of \mathcal{K}_1^α is exactly the first Betti number of K .*

Proof: In general there are many different S-reduction sequences which achieve the goal of the theorem, and we present only one particular example. In a first step, one can remove all 2-cells of \mathcal{K} via a sequence of elementary reductions. To see this, note that as a consequence of the Jordan curve theorem, any planar CW-complex must have a 2-cell with a free edge — one just has to pick a 2-cell at the outer boundary of the complex. Applying an elementary reduction not only gives a reduced S-complex, but in fact a reduced planar CW-complex which is still connected. By applying this argument successively, one 2-cell after another can be removed from the planar CW-complex, without leaving the category of planar CW-complexes which are connected.

After the first step, one is left with a connected planar CW-complex of dimension one. The addition of the empty set to the associated chain complex in dimension -1 then allows one to perform a first coreduction, which removes a vertex from the S-complex. From this point onwards, one is no longer working with CW-complexes, but rather with S-complexes. Using a succession of edge-vertex coreductions it is now possible to remove all the remaining vertices of the S-complex, as well as all the edges in a spanning tree for the 1-skeleton of K , i.e., the resulting α -reduced S-complex \mathcal{K}^α has the structure stated in the theorem. Finally, since S-reductions preserve the homology of the chain complex and since the final S-complex has a vanishing boundary map, the

¹It is straightforward to see that every CW-complex has to contain at least one 0-cell.




Figure 6: This example shows that in general it is not possible to completely reduce a two-dimensional non-planar CW-complex using elementary coreductions or elementary reductions. The image shows a CW-complex representing a torus, where vertices with identical letters are identified, as well as edges between them. In order to start the coreduction sequence we add the empty set in dimension -1 .

number of remaining 1-cells equals the first Betti number of the original CW-complex. This completes the proof of the theorem. \diamond

In fact, the following theorem follows easily from our above discussion.

Theorem 3.7 *Let K be a planar regular CW-complex with a CW-complex structure \mathcal{K} . Then the homology of K can be computed in $O(|\mathcal{K}|)$ time. For this, the Smith normal form algorithm does not have to be employed and no knowledge of the incidence indices of cells of \mathcal{K} is necessary.*

Let us remark that the results of Theorems 3.6 and 3.7 may be extended to the case of homology generators. To see this observe that in the setting of Theorem 3.6 the set of generators of the α -reduction coincides with the set of the homology generators of the α -reduction. Therefore, to obtain the set of generators of the original S-complex (S, κ) it is enough to find the images of these generators in the map ι^α given by (2). By [21, Theorem 3.1] the cost of finding this image is $O(|\mathcal{K}|w(\mathcal{K}))$. Therefore we have the following theorem.

Theorem 3.8 *Let K be a planar regular CW-complex with a CW-complex structure \mathcal{K} . Then the time complexity of computing the generators of the homology of K is $O(|\mathcal{K}|w(\mathcal{K})g(\mathcal{K}))$, where $g(\mathcal{K})$ denotes the number of homology generators of K .*

In particular, in the setting of two-dimensional raster images the cost is $O(|\mathcal{K}|g(\mathcal{K}))$.

Before closing this section, we would like to point out that the conclusion of Theorem 3.6 and Theorem 3.7 are in general not true for two-dimensional non-planar CW-complexes. To see this, consider the torus represented by the CW-complex shown

in the left image of Figure 6. As usual, vertices with the same letter are identified, as well as the corresponding edges between them. After appending the empty set in dimension -1 as before, one can remove the center vertex of the CW-complex. Subsequent edge-vertex coreductions then allow one to remove all the remaining vertices, as well as all the edges in a spanning tree of the 1-skeleton, resulting in the S-complex shown in the center image of Figure 6. Finally, a sequence of 2-cell-edge coreductions can be used to remove nine of the 2-cells, and this furnishes the S-complex on the right. Notice that for this S-complex, no further elementary reductions or coreductions are possible.

3.3 Rectangular CW-Complexes

The previous section identified a situation in which the coreduction algorithm always leads to a final S-complex with trivial boundary operators. In such a case, the homology of the underlying CW-complex can be determined directly from the generators of the remaining chain groups, and no application of the Smith normal form algorithm is necessary at all. In most situations, however, the final S-complex after the coreduction step does contain nontrivial boundary operators. In order to then efficiently apply the Smith normal form algorithm, one has to be able to determine the matrix representations for the boundary operators quickly, and this in turn depends on the efficient computability of the incidence index κ . Since the reduced S-complex is no longer a regular CW-complex, we cannot use the equations in (c) and (d) of Theorem 3.1. We demonstrated this in Example 3.5. We can always use the algorithm in Table 1 to obtain the incidence index κ before we start the reductions and use them after the reductions to build the matrices of the boundary map in the remaining complex. It will be shown in this section that CW-complexes in a special class, called rectangular CW-complexes, are effective S-complexes. Therefore, in this case the computation of all incidence numbers before the S-reductions start is not needed, because the incidence index κ can easily be computed directly from the representation of the cells, even though the S-complex obtained after the reduction process is no longer a rectangular cell complex. The explicit formula for the incidence of two cells which can be used for the remaining cells in the reduced complex is provided in this section.

In the following, the term *interval* is always used for a compact interval $I = [a, b]$ in \mathbb{R} with $a \leq b$. We say that the interval I is *degenerate* if $a = b$, otherwise it is called *nondegenerate*. A *facet* of an interval is defined as follows: If the interval I is degenerate, then it has no facets; if I is nondegenerate, then it has two facets, which are given by the degenerate intervals $[a] := [a, a]$ and $[b] := [b, b]$. The class of regular CW-complexes which is studied in this section is formed via rectangles. For this, a *rectangle* in \mathbb{R}^d is a product $Q = Q_1 \times Q_2 \times \cdots \times Q_d$ of d intervals Q_1, Q_2, \dots, Q_d . The *dimension* of Q , which is denoted by $\dim Q$, equals the number of nondegenerate intervals in the sequence Q_1, Q_2, \dots, Q_d . Next, we need to define the notion of facets

for rectangles. For this, assume $d > 1$, and let $P = P_1 \times P_2 \times \cdots \times P_d$ denote a second rectangle in \mathbb{R}^d . Then we say that P is a *facet* of Q , which is denoted by $P < Q$, if $P \subset Q$, the dimensions satisfy $\dim P = \dim Q - 1$, and there exists a j such that P_j is a facet of Q_j . It can easily be seen that in this case the index j is unique and that $\dim P_i = \dim Q_i$ for all $i \neq j$. Finally, P is called a *face* of Q , if either there is a descending sequence of facets joining Q to P , or $P \in \{\emptyset, Q\}$. The round of preparatory definitions is completed by the notion of boundary. For an interval $[a, b]$ we define its *boundary* as

$$\text{bdr } [a, b] := \begin{cases} \{a, b\} & \text{if } a \neq b, \\ \emptyset & \text{if } a = b. \end{cases}$$

Now let $d > 1$ and consider a rectangle $Q = Q_1 \times Q_2 \times \cdots \times Q_d$. Then its *boundary* is defined as $\text{bdr } Q := \bigcup_{i=1}^d Q_1 \times \cdots \times \text{bdr } Q_i \times \cdots \times Q_d$. It can easily be verified that if one thinks of an n -dimensional rectangle Q as an n -dimensional manifold with boundary which is embedded in \mathbb{R}^d , then the set $\text{bdr } Q$ contains exactly the boundary points of the manifold.

After the preparations of the previous paragraph, we can now define the central object for this section. A *rectangular structure* is a finite collection \mathcal{Q} of rectangles such that for any choice of $P, Q \in \mathcal{Q}$ one has either $P \cap Q = \emptyset$, or $P \cap Q$ is a common face of both P and Q which in addition belongs to \mathcal{Q} . Moreover, we assume that for any rectangle $Q \in \mathcal{Q}$, its boundary satisfies $\text{bdr } Q = \bigcup \{P \in \mathcal{Q} : P \text{ is facet of } Q\}$. Finally, a *rectangular CW-complex* is given by the union of some rectangular structure, i.e., it is the subset of Euclidean space which is occupied by the rectangles in \mathcal{Q} . Any 0-dimensional rectangle $Q \in \mathcal{Q}$ is called a *vertex*, and if $Q \in \mathcal{Q}$ is an n -dimensional rectangle, then the set $Q \setminus \text{bdr } Q$ is called an *n -cell*. For rectangular CW-complexes, we have the following theorem.

Theorem 3.9 *Any rectangular CW-complex is a regular CW-complex.*

Proof: Any rectangular CW-complex is the finite union of closed rectangles, and therefore it is a compact subset of \mathbb{R}^d with respect to the standard topology. In particular, it is a Hausdorff space.

Let \mathcal{Q} denote the rectangular structure that induces the rectangular CW-complex and let \mathcal{Q}_n denote the set of n -dimensional rectangles in \mathcal{Q} . Then $K^n := \bigcup \{Q \in \mathcal{Q} : \dim Q \leq n\}$ is an n -dimensional skeleton of the rectangular CW-complex. Obviously, $\mathcal{Q}_n \subset \mathcal{Q}$ is finite for all $n = 0, 1, \dots, d$.

Now consider a rectangle $Q = Q_1 \times Q_2 \times \cdots \times Q_d \in \mathcal{Q}_n$. It is well-known that there exists a homeomorphism $\hat{h}_n : B_1^n(0) \rightarrow (0, 1)^n$ that can be extended to a homeomorphism $h_n : \overline{B_1^n(0)} \rightarrow [0, 1]^n$. Thus, it suffices to construct a homeomorphism $f_Q : [0, 1]^n \rightarrow Q$ which attaches Q to K^{n-1} . Due to $\dim Q = n$, there exists $\{i_1, \dots, i_n\} \subset \{1, \dots, d\}$ such that Q_{i_1}, \dots, Q_{i_n} are the nondegenerate intervals in Q . For $i \notin \{i_1, \dots, i_n\}$ we write $Q_i = [q_i]$, and for $j \in \{i_1, \dots, i_n\}$ we let $Q_j = [a_j, b_j]$,

where $a_j \neq b_j$. Using this notation, consider the map $f_Q : [0, 1]^n \rightarrow Q$ which is defined on points $(x_1, \dots, x_n) \in [0, 1]^n$ by

$$(f_Q(x_1, \dots, x_n))_j := \begin{cases} a_j + x_i(b_j - a_j) & \text{for } j = i_i \in \{i_1, \dots, i_n\}, \\ q_j & \text{for } j \notin \{i_1, \dots, i_n\}, \end{cases}$$

where $(f_Q(x_1, \dots, x_n))_j$ denotes the j -th component of the vector $f_Q(x_1, \dots, x_n)$. One can readily see that the inverse map $f_Q^{-1} : Q \rightarrow [0, 1]^n$ is given by

$$(f_Q^{-1}(q_1, \dots, q_d))_j = \frac{q_{i_j} - a_{i_j}}{b_{i_j} - a_{i_j}} \in [0, 1]^n,$$

for $j = 1, \dots, n$ and $(q_1, \dots, q_d) \in Q$. Both mappings f_Q and f_Q^{-1} are continuous bijections, and therefore $f_Q : [0, 1]^n \rightarrow Q$ is a homeomorphism for any $Q \in \mathcal{Q}$. The fact that $f_Q(\partial[0, 1]^n) = \text{bdr } Q \subset K^{n-1}$ follows readily from the definition of f_Q and the assumption that $\text{bdr } Q = \bigcup\{P \in \mathcal{Q} : P \text{ is facet of } Q\}$. \diamond

The next result shows that for rectangular CW-complexes, the incidence index κ can be computed easily and efficiently.

Theorem 3.10 *Consider a rectangular CW-complex given by a rectangular structure \mathcal{Q} . Let P and Q denote two arbitrary rectangles in \mathcal{Q} with $\dim Q = 1 + \dim P$, and define the number α_{QP} as follows. For $d = 1$ and $Q = [a, b]$ let*

$$\alpha_{QP} := \begin{cases} -1 & \text{if } P = [a], \\ 1 & \text{if } P = [b], \\ 0 & \text{otherwise,} \end{cases}$$

and for $d > 1$ set

$$\alpha_{QP} := \begin{cases} (-1)^{\sum_{i=1}^{j-1} \dim Q_i} \alpha_{Q_j P_j} & \text{if } P < Q \text{ and } j \text{ satisfies } P_j < Q_j, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Then the numbers α_{QP} are incidence indices for the given rectangular CW-complex, i.e., they satisfy all the conditions of Theorem 3.1.

Proof: The above definition of α_{QP} immediately implies that (a) and (b) of Theorem 3.1 are satisfied. In addition, condition (c) is clear in the case $d = 1$, and it follows easily also for $d > 1$, since the index j and the corresponding exponent

$$\varepsilon(Q, P) := \sum_{i=1}^{j-1} \dim Q_i$$




Figure 7: Illustrations of Case 1 (left) and Case 2 (right) in the proof of Theorem 3.10.

in (4) are the same for the two vertices of a 1-cell. Thus, we only have to establish the validity of condition (d) in Theorem 3.1. For this, consider $d > 1$ and let Q , P , R , and S denote rectangles with dimensions n , $n - 1$, $n - 1$, and $n - 2$, respectively. Furthermore, assume that P and R are facets of Q , and that S is the unique facet of both P and R . Finally, let Q_i , P_i , R_i , and S_i denote the i -th component interval of Q , P , R , and S , respectively. Using this setting, we have to verify the identity

$$\alpha_{QP} \alpha_{PS} + \alpha_{QR} \alpha_{RS} = 0 \quad (5)$$

in order to establish assumption (d) in Theorem 3.1. For this, we distinguish the following two cases.

Case 1 (see Figure 7 left): Assume that the index j in the definition of a facet is the same for both $R < Q$ and $P < Q$.

Under this additional assumption, one obviously has $\varepsilon(Q, P) = \varepsilon(Q, R)$. Moreover, since $P \cap R \supset S \neq \emptyset$, we see that $P_j = R_j = S_j$ and they coincide with one of the endpoints of Q_j . This results in $\alpha_{QP} = \alpha_{QR}$, and therefore it remains to be shown that we have

$$\alpha_{PS} + \alpha_{RS} = 0 .$$

Since $S < P$ and $S < R$, there exist unique integers k and l such that $S_k < P_k$ and $S_l < R_l$. In fact, we must have $k = l$, since otherwise one would obtain $\dim S = n - 3$, which contradicts our hypothesis. Thus, one can readily see that $\varepsilon(P, S) = \varepsilon(R, S)$. In addition, we have $S_k = P_k \cap R_k = [b]$ for some real number b . Without loss of generality, $P_k = [a, b]$ and $R_k = [b, c]$, and therefore

$$\alpha_{PS} = (-1)^{\varepsilon(P,S)} \alpha_{P_k S_k} = (-1)^{\varepsilon(P,S)} , \quad \alpha_{RS} = (-1)^{\varepsilon(R,S)} \alpha_{R_k S_k} = (-1)^{\varepsilon(P,S)+1} .$$

From this the conclusion follows.

Case 2 (see Figure 7 right): *There are two distinct indices $j \neq k$ such that $P_k < Q_k$ and $R_j < Q_j$.*

Under this assumption, the intervals P_k and R_j are degenerate, and all other interval components of P and R with coordinates i for which $\dim Q_i = 1$ are nondegenerate. Without loss of generality, we may assume $k < j$. The degenerate interval P_k may be either a left or a right endpoint of the nondegenerate interval Q_k , and R_j may be either a left or a right endpoint of the nondegenerate interval Q_j . This gives rise to four distinct possibilities of which we only describe one in detail; the remaining cases can be treated analogously and are left to the reader. Consider the 2-dimensional rectangle in the $k \times j$ coordinates given by

$$Q' := Q_k \times Q_j = [a, b] \times [c, d] ,$$

as well as the edge parts

$$\begin{aligned} P' &:= P_k \times P_j = [a] \times [c', d'] , & \text{with } c \leq c' < d' \leq d , \\ R' &:= R_k \times R_j = [a', b'] \times [c] , & \text{with } a \leq a' < b' \leq b . \end{aligned}$$

Due to $P \cap R \supset S \neq \emptyset$, we then must have

$$S' := S_k \times S_j = [a] \times [c] .$$

This in turn is only possible if $c' = c$ and $a' = a$, i.e., the interval P' intersects the interval R' in the lower-left corner of Q' . (Notice that the above-mentioned remaining three cases correspond to the remaining three corners of Q' .) The definitions

$$p = \varepsilon(Q, P) = \sum_{i=1}^{k-1} \dim Q_i \quad \text{and} \quad r = \varepsilon(Q, R) = \sum_{i=1}^{j-1} \dim Q_i$$

then furnish

$$\alpha_{QP} = (-1)^p \alpha_{Q_k P_k} = (-1)^{p+1} \quad \text{and} \quad \alpha_{QR} = (-1)^r \alpha_{Q_j R_j} = (-1)^{r+1} .$$

According to $\dim P_k = \dim S_k = 0$, the sum in the formula for the exponent $\varepsilon(P, S)$ extends up to the index $j - 1$, analogously to the case of the exponent r defined above. In other words, we have

$$\varepsilon(P, S) = \sum_{i=1}^{j-1} \dim P_i .$$

Now observe that $\dim P_i = \dim Q_i$ for all $i \leq j - 1$ — except for the choice $i = k$, which corresponds to $\dim Q_k = 1$ and $\dim P_k = 0$. Hence, the identity $\varepsilon(P, S) = r - 1$ holds and we obtain

$$\alpha_{PS} = (-1)^{r-1} \alpha_{P_j S_j} = (-1)^r .$$

In addition, we have $\dim R_i = \dim Q_i$ for all $i \leq j - 1$. Combined with the assumption $k < j$ and $S_k < R_k$, this furnishes

$$\varepsilon(R, S) = \sum_{i=1}^{k-1} \dim R_i = \sum_{i=1}^{k-1} \dim Q_i = p ,$$

as well as

$$\alpha_{RS} = (-1)^p \alpha_{R_k S_k} = (-1)^{p+1} .$$

This finally implies

$$\alpha_{QP} \alpha_{PS} + \alpha_{QR} \alpha_{RS} = (-1)^{p+1} (-1)^r + (-1)^{r+1} (-1)^{p+1} = 0 ,$$

which completes the proof of the theorem. \diamond

Remark 3.11 The formula (4) is inspired by the boundary operator formula for cubical sets which was derived in [13, Corollary 2.35]. In fact, an alternative presentation of rectangular CW-complexes can be given in terms of rescaling of cubical sets as defined in [13, Section 6.4.2]. Suppose that \mathcal{Q} is a rectangular structure for $X \subset \mathbb{R}^d$ such that every cube $Q \in \mathcal{Q}$ is the product of intervals with rational coordinates. This hypothesis is not restrictive in practice, because the interval arithmetic always produce such intervals. Let α_i be a common multiple of all denominators present in the i -th coordinate intervals of all these cubes. Consider the scaling isomorphism $\Lambda^\alpha : \mathbb{R}^d \rightarrow \mathbb{R}^d$ defined by $\Lambda^\alpha(x) := (\alpha_1 x_1, \alpha_2 x_2, \dots, \alpha_d x_d)$. Then the image X^α of X under the mapping Λ^α is a cubical set. Moreover, it can be seen that the chain complex associated with X is isomorphic to the cubical subcomplex of $\mathcal{C}(X^\alpha)$ whose n -th chain group is generated by the chains

$$\hat{Q} := \sum \left\{ \hat{P} : P \in \mathcal{K}_n(Q^\alpha) \right\}$$

over $Q \in \mathcal{Q}$, where $\mathcal{K}_n(Q^\alpha)$ denotes the set of elementary n -cubes in Q^α and \hat{P} is the dual generator corresponding to $P \in \mathcal{K}_n(X^\alpha)$. We do not present the proof of this fact because it is not used in the computations of this paper.

4 Numerical Experiments

In this final section of the paper we present a first application of the coreduction method for regular CW-complexes in the context of planar rectangular CW-complexes. For this, we consider again the study of topological properties of nodal domains of functions, which has been considered before in [1, 4, 5, 10, 16, 17]. In particular, we will improve on the numerical studies presented in [5]. As was already mentioned in the introduction, the latter study developed a rigorous computational technique for finding the homology of nodal sets of smooth functions. In order to determine and rigorously validate the

boundary of nodal sets in the plane, the method of [5] finds a non-uniform cubical decomposition of the underlying two-dimensional domain of the function in such a way that the signs of the function values on the corners of any square in the decomposition completely describe the structure of the nodal line within this square: If all signs are the same, the nodal line has an empty intersection with the square; if exactly one corner has a different sign from the remaining three corners, then the nodal line is a simple curve which originates/ends at the two edges adjacent to the exceptional corner, etc. Once the algorithm has determined a non-uniform decomposition of this type, the only way to compute the homology of the nodal sets using homology codes which were available at the time was to embed the non-uniform grid into a fine uniform grid, whose size is determined by the smallest square in the non-uniform decomposition. It was already pointed out in [5], that this approach severely limits the practical applicability of the method.

In order to address these shortcomings with respect to the homology computation in [5] we have implemented a coreduction algorithm for planar rectangular CW-complexes. According to Theorem 3.6 and Theorem 3.7, in this particular situation it is possible to employ elementary reductions and coreductions to completely reduce the underlying S-complex, i.e., at the end of the algorithm all boundary operators are trivial and there is no need for an application of the Smith normal form algorithm. In the remainder of this section, we describe the improvements that the coreduction method brings over the homology computation method employed in [5]. As mentioned above, in [5] the homology of the nodal domains was computed using a uniform cubical approximation, whose grid size was determined by the smallest cube in the non-uniform adaptive grid, and whose homology was then computed using the standard coreduction algorithm for uniform cubical complexes developed in [19]. While the latter algorithm is extremely efficient, the insistence on uniform cubical grids puts severe restrictions on the largest possible problem size that could be tackled in [5]. For the case of nodal domains of doubly-periodic random trigonometric polynomials, which will be described in more detail below, it was only possible to consider fine uniform grids of sizes up to 65000^2 . In order to compare the coreduction algorithm for rectangular non-uniform CW-complexes to the one for uniform cubical complexes, we consider the test case of random trigonometric polynomials. In previous theoretical work [16], rigorous bounds on the probability of a correct homology computation using uniform cubical grids of certain sizes were derived — and the validity of these estimates as a function of the grid size for random trigonometric polynomials of certain degrees was established in [5]. These studies are concerned with a special class of random periodic Gaussian fields, namely, with random trigonometric polynomials on the base domain $\Gamma = [0, 1]^2$ of the




Figure 8: Sample nodal domain patterns for random trigonometric polynomials in two dimensions. From left to right the images correspond to the degrees $K = 16, 20, 24$ in (6). In each image, a yellow outline is selected, whose blow-up is shown in the second row. These blow-ups contain the adaptive numerical grid which is used for homology verification.

form

$$\begin{aligned}
 f(x_1, x_2, \omega) = \sum_{k,\ell=0}^K \alpha_k \alpha_\ell \cdot & (g_{k,\ell,1}(\omega) \cos(2\pi k x_1) \cos(2\pi \ell x_2) \\
 & + g_{k,\ell,2}(\omega) \cos(2\pi k x_1) \sin(2\pi \ell x_2) \\
 & + g_{k,\ell,3}(\omega) \sin(2\pi k x_1) \cos(2\pi \ell x_2) \\
 & + g_{k,\ell,4}(\omega) \sin(2\pi k x_1) \sin(2\pi \ell x_2)) .
 \end{aligned} \tag{6}$$

In (6), the degree of the trigonometric polynomial is given by $K \geq 3$, and the factors $g_{k,\ell,m}$ are random variables defined over a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which are independent and normally distributed with mean 0 and variance 1. See also [4, 17].

Studying the nodal domains of random trigonometric polynomials of the form (6) provides an excellent testing ground for the new coreduction algorithm for non-uniform planar rectangular CW-complexes. This is due to the fact that with increasing degree K both the topological complexity of the resulting nodal domains, and the complexity of the non-uniform adaptive grid produced by the algorithm of [5] grows. In fact, our




Figure 9: The left image shows the growth of the averaged 0-th Betti numbers (blue curve) and 1-st Betti numbers (red curve) of the nodal domains of random trigonometric polynomials given by (6); both axes in this image use a logarithmic scale. The image on the right shows the average number of validated rectangles in the final adaptive grid generated by the validation algorithm in [5] as a function of the degree K of the random trigonometric polynomial; in this image, only the vertical axis is using logarithmic scaling.

simulations — which are described in more detail below — show that the averaged Betti numbers $\beta_0(K)$ and $\beta_1(K)$ of random trigonometric polynomials of degree K roughly follow the proportionalities

$$\beta_0(K) \sim K^{1.337} \quad \text{and} \quad \beta_1(K) \sim K^{2.272} .$$

Furthermore, the size of the adaptive validated grid grows exponentially via

$$\text{adaptive grid size} \sim 1.306^K .$$

Thus, for the largest simulations that are included in the discussion below, we are dealing with non-uniform rectangular CW-complexes with millions of cells. Nevertheless, even in these extreme cases the runtime of our homology algorithm turns out to be only a few seconds.

The graphs shown in Figures 10 and 11 contain basic timings of the coreduction algorithm for non-uniform planar rectangular CW-complexes, as well as comparisons with the method employed in [5]. To obtain these graphs, we applied the validation algorithm of [5] to random ensembles of trigonometric polynomials of the form (6), for values of the degree K between 3 and 32. The sizes of the random ensembles varied between 3000 for small values K and 5 for the largest few values. The selection of the ensemble sizes was dictated by the running time of the validation algorithm. Once the non-uniform validated adaptive grid has been determined, we applied the coreduction algorithm for non-uniform planar rectangular CW-complexes to the non-uniform grid

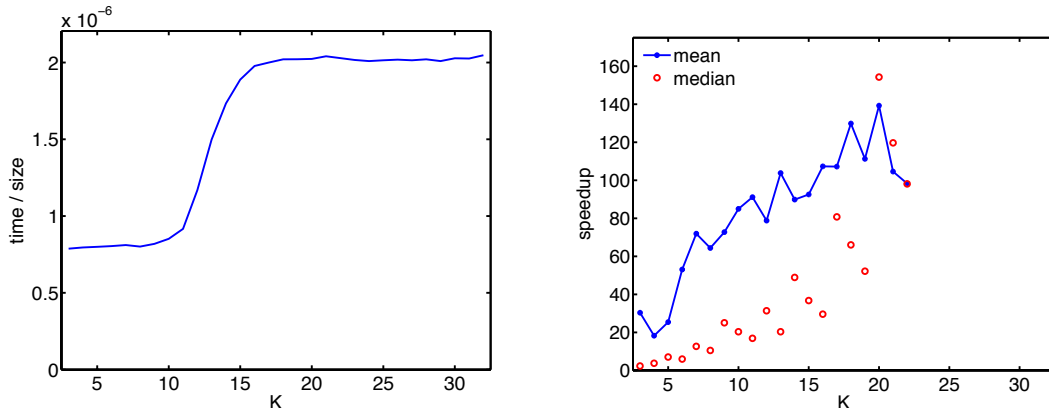


Figure 10: Timings for the coreduction algorithm for planar rectangular CW-complexes. The left image shows the ratio of computational time over the actual size of the considered non-uniform planar rectangular CW-complex, averaged over the samples of a random ensemble of trigonometric polynomials, as a function of the degree K . This image indicates linear time complexity, except for the increase in computational time around $K \approx 15$ which is caused by memory paging effects. The right image shows the average and median speedups which were observed in comparison with the uniform grid method employed in [5]. Notice that no comparisons could be made for $K \geq 23$.

to determine the Betti numbers. In addition, whenever possible — i.e., if the size $1/M$ of the smallest cube in the adaptive grid was not too small — we applied the homology algorithm used in [5].

The first set of timings obtained in this way can be found in the left image of Figure 10. This image shows the ratio of computational time of the coreduction algorithm for non-uniform planar rectangular CW-complexes divided by the actual size of the considered non-uniform rectangular CW-complex, averaged over the samples of a random ensemble of trigonometric polynomials, as a function of the degree K . This image indicates that the runtime complexity of the algorithm is linear in the complex size, as indicated by the almost constant ratios as a function of K . At first glance, the sudden increase of this ratio around $K \approx 15$ seems to contradict the linearity. However, for degrees less than this threshold the resulting complexes are so small that they can be completely dealt with without memory paging. Once the degree of the random trigonometric polynomial exceeds the threshold value, memory paging has to be employed by the C++ program to deal with the larger complex size — and this results in the increased computational effort.

In order to compare the new algorithm with the homology algorithm used in [5], we also computed the mean and median speedups for the runtimes of the two algorithms. The results of these computations can be found in the right image of Figure 10. The data shows that significant average speedups of two orders of magnitude can be observed for degrees K larger than 10, while the speedups are more moderate for smaller




Figure 11: Further comparisons of the coreduction algorithm for planar rectangular CW-complexes with the uniform grid method employed in [5]. The top left image shows the percentage of samples for which the algorithm in [5] could not compute the homology due to memory constraints. The top right image gives the maximal speedup which was observed within the random ensembles for each value of the degree K (notice the logarithmic scale on the vertical axis). The decay of the maximal speedup for K -values approaching 23 is due to the fact that the algorithm in [5] is able to compute the homology of nodal domains of trigonometric polynomial of such degrees only if the topology of the nodal domain is very simple. In these cases, one cannot see the huge advantage of the new approach when compared with the uniform grid method. The lower left image shows the distribution of $\log_2 M$, where $1/M$ is the size of the smallest cube in the adaptive grid, for various degrees K , and the lower right image presents the averaged speedups as a function of $\log_2 M$, for K -values between 3 and 21.

values of K — due to the low complexity of the nodal domains. We would like to point out, however, that for $K \geq 23$ no speedup data could be obtained. An explanation of this fact is given in the top left image of Figure 11. As a function of K , this image depicts the percentage of samples within each ensemble of random trigonometric polynomials with fixed degree K , for which it was impossible to determine the homology using the method of [5]. In fact, for none of the considered functions of degrees $K \geq 23$ the old method could be applied, since the size $1/M$ of the smallest cube in the validated adaptive grid was many orders of magnitudes too small. This failure was the reason to restrict the studies in [5] to $K \leq 16$. Nevertheless, in all of these cases, the new algorithm could produce the Betti numbers in no time.

As a final measure of the performance of the two algorithms we show the maximal observed speedup within each ensemble of random trigonometric polynomials with fixed degree K in the top right image of Figure 11. These maximal speedups were about four orders of magnitude for smaller values of K , and decrease to two orders of magnitude as K increases from 10 to 22. Notice, however, that the latter numbers do not give an accurate picture. This is due to the fact that the range of K -values where the maximal speedups decrease corresponds exactly to the range of K -values for which the failure rate of the old homology code increases. In other words, for these degrees, comparisons can only be made if the nodal domain geometry is simple enough to allow for the application of the homology algorithm of [5] — and these of course are exactly the cases where we do not expect large speedups from the new code. In fact, this statement is illustrated more in the bottom pictures of Figure 11. In these images, we take into account the size M^2 of the homogeneous grid which is used for the homology computations in [5], where again $1/M$ denotes the side length of the smallest cube in the validated adaptive grid. The image in the lower left shows the distribution of $\log_2 M$ for various degrees K , which indicates a fairly narrow range for each K , which of course shifts to the right as K increases. In the lower right image we then present the averaged speedups as a function of $\log_2 M$, for K -values between 3 and 21. This image clearly shows that the decay in the upper right image of Figure 11 is due to the lack of comparison data from the old algorithm.

Acknowledgments

T.K. was supported by the Discovery Grant from NSERC of Canada. M.M. was partially supported by MNiSW grant N N201 419639. T.W. was partially supported by NSF grants DMS-0639300 and DMS-0907818, as well as the U.S. Department of Energy under Contract DE-FG02-05ER25712.

References

- [1] R. J. Adler and J. E. Taylor. *Random Fields and Geometry*. Springer-Verlag, New York, 2007.
- [2] D. Blömker, S. Maier-Paape, and T. Wanner. Phase separation in stochastic Cahn-Hilliard models. In A. Miranville, editor, *Mathematical Methods and Models in Phase Transitions*, pages 1–41. Nova Science Publishers, New York, 2005.
- [3] G. E. Bredon. *Topology and Geometry*, volume 139 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [4] S. Day, W. D. Kalies, K. Mischaikow, and T. Wanner. Probabilistic and numerical validation of homology computations for nodal domains. *Electronic Research Announcements of the American Mathematical Society*, 13:60–73, 2007.
- [5] S. Day, W. D. Kalies, and T. Wanner. Verified homology computations for nodal domains. *SIAM Journal on Multiscale Modeling & Simulation*, 7(4):1695–1726, 2009.
- [6] C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for betti numbers of simplicial complexes. In *SCG '93: Proceedings of the ninth annual symposium on Computational geometry*, pages 232–239, New York, NY, USA, 1993. ACM.
- [7] J.-G. Dumas, F. Heckenbach, D. Saunders, and V. Welker. Computing simplicial homology based on efficient Smith normal form algorithms. In *Algebra, geometry, and software systems*, pages 177–206. Springer, Berlin, 2003.
- [8] H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.
- [9] J. Friedman. Computing Betti numbers via combinatorial Laplacians. *Algorithmica*, 21(4):331–346, 1998.
- [10] M. Gameiro, K. Mischaikow, and T. Wanner. Evolution of pattern complexity in the Cahn-Hilliard theory of phase separation. *Acta Materialia*, 53(3):693–704, 2005.
- [11] M. Juda and M. Mrozek. \mathbb{Z}_2 -homology of weak 2-pseudomanifolds may be computed in $O(n\alpha(n))$ time. Preprint, Jagiellonian University, Computer Science Department, 2010. Available online at <http://www.ii.uj.edu.pl/~mrozek/papers/homology-of-2-manifolds-15.pdf>.
- [12] T. Kaczynski, K. Mischaikow, and M. Mrozek. Computing homology. *Homology, Homotopy and Applications*, 5(2):233–256, 2003.
- [13] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational Homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004.
- [14] T. Kaczyński, M. Mrozek, and M. Ślusarek. Homology computation by reduction of chain complexes. *Computers & Mathematics with Applications*, 35(4):59–70, 1998.
- [15] W. S. Massey. *A Basic Course in Algebraic Topology*, volume 127 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [16] K. Mischaikow and T. Wanner. Probabilistic validation of homology computations for nodal domains. *Annals of Applied Probability*, 17(3):980–1018, 2007.
- [17] K. Mischaikow and T. Wanner. Topology-guided sampling of nonhomogeneous random processes. *Annals of Applied Probability*, 20(3):1068–1097, 2010.
- [18] M. Mrozek. Cech type approach to computing homology of maps. *Discrete and Computational Geometry*, 2010. In print.

- [19] M. Mrozek and B. Batko. Coreduction homology algorithm. *Discrete & Computational Geometry*, 41(1):96–118, 2009.
- [20] M. Mrozek, P. Pilarczyk, and N. Żelazna. Homology algorithm based on acyclic subspace. *Computers & Mathematics with Applications*, 55(11):2395–2412, 2008.
- [21] M. Mrozek and T. Wanner. Coreduction homology algorithm for inclusions and persistent homology. *Computers & Mathematics with Applications*, to appear.
- [22] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Menlo Park, 1984.
- [23] S. Peltier, A. Ion, W. Kropatsch, G. Damiand, and Y. Haxhimusa. Directly computing the generators of image homology using graph pyramids. *Image and Vision Computing*, 27(7):846–853, 2009.
- [24] A. Storjohann. Near optimal algorithms for computing smith normal forms of integer matrices. In *Proceedings of the 1996 international symposium on Symbolic and algebraic computation, ISSAC '96*, pages 267–274, New York, NY, USA, 1996. ACM.
- [25] T. Wanner, E. R. Fuller Jr., and D. M. Saylor. Homological characterization of microstructure response fields in polycrystals. *Acta Materialia*, 58(1):102–110, 2010.