JOURNAL OF SPATIAL INFORMATION SCIENCE
Article submitted for review

Hierarchical Euler-Poincaré operators

Patrick Erik Bradley

Karlsruher Institut für Technologie, Institut für Photogrammetrie und Fernerkundung,
Englerstr. 7, 76131 Karlsruhe, Germany

October 25, 2010

Abstract: Motivated by segmentation and cartographic segmentation, a relational model for hierarchies of topologies is introduced together with their Euler-Poincaré operators in N dimensions. Topological consistency leads to conditions for computing boundary operators from the highest level of detail to the coarser levels in the hierarchy, which can be interpreted in terms of ultrametric integration.

Keywords: topology, Euler operators, generalisation, ultrametricity, chain complex, persistence complex, relational model, N-dimensional spatial modelling

1 Introduction

Topology has undoubtedly become an issue in geographic or other spatial information sciences and systems. It abstracts from geometry and allows to model neighbourhood relations without having to provide for a distance. Topological properties of a 'space' allow a first coarse classification of such entities. One of the first methods using concepts of set-theoretic topology for geographic information can be found in [5]. Like Max Egenhofer recently did, we also advocate that 'space' need not refer to the conventional 3D physical space only: any entity which a mathematician or physicist would call 'space' can be taken into consideration, as long as it has some underlying topology[1]. Usually, one considers 'spaces' which are somehow related to the physical or geographic space we encounter every day, and many of these spaces have a geometry or some other rich structure on top of their topology. In particular, they can in

[1]Probably the most general concept would be an object in some appropriate category.

principle be of any finite or infinte dimension[2]. Thinking, for example, of a gray-scale image, we obtain a real-valued function defined on a grid, and the Fourier transform expresses this function in terms of its harmonics at all different frequencies. This transform is therefore an operation on an infinite-dimensional space of functions. In fact, the motivation for this article comes from image processing, where generalisation, in various forms of analytic methods, plays a crucial role in the finding of features or of homogeneous segments. The following section will treat this aspect from a topological point of view.

Albeit its fundamental importance for spatial information, the treatment of topology often does not go beyond the level of graphs or neighbourhood relations. Application of the deep theory of algebraic topology seems merely at its initial stage, due to a lack of interdiscipliary efforts between topologists and spatial information modellers, or 'topological modellers' as we will call the latter. As evidence for this we refer to the fact that already the step from 2D to 3D has posed a serious problem to the topological modeller. Whereas the planar graph was found of use for constructing a relational database model for two dimensional spatial data, relational 3D databases were often defined as complicated graph models in order to cope with each extra dimension. Only recently did simplicial or cellular complexes find their way into spatial information systems for three or four dimensions [4]. The boundary representation method from volume modelling has also been taken into consideration. Although it is of very topological nature, the modelling of an object by its boundary is problematic in that it cannot distinguish between a 'filled' and an 'unfilled' boundary cycle, i.e. it misses entirely the topologically essential notion of 'hole' as being a cycle which is not itself the boundary of a cell. A model suitable for any finite dimension is provided by G-maps [7]. However, these are of fairly high storage complexity [2].

The foundational paper [3] emphasises the structural independence of dimension in the chain complex, the topologist's favourite topological model of a space, and constructs a relational version thereof for the use by the topological modeller. Crucial here is the *boundary operator*. Instead of representing a cell by its boundary, as in B-Rep, the topological point of view takes all cells into its model and then implements the boundary relation by saying that any given cell has at its boundary all cells with some given orientation and multiplicity, and fullfills the algebraic form of the consistency relation which says that the boundary itself has no boundary. The boundary operator is thus represented by a sparse matrix ∂ such that $\partial^2 = 0$, and in [3] it is observed that this idea has a straightforward relational formulation using partial matrices. This construction will be reviewed in more detail in Section 3, together with the corresponding generalised Euler operators known in volume modelling [8]. The latter respect the *Euler-Poincaré equation* for chain complexes, a generalisation of the known Euler formula of polytopes to N dimensions.

Section 4 extends the reviewed notions from Section 3 to the hierarchical situation as it can be encountered in generalisation. Here, the categorial point of view taken

[2]However that notion may be defined!

in [3] comes into play, as the generalisation maps between the coarser and the finer levels of detail are instances of a relative situation given by *morphisms*. Here, these are given by continuous maps in one direction and (relational) chain maps in the other. Sequences of Euler operators act on the resulting spaces of so-called *hierarchical relational chain complexes*. In this way, the *hierarchical Euler-Poincaré operators* can be endowed with a norm which can be interpreted as a range of its action.

Section 5 makes the final crucial observation that hierarchical relational chain complexes have an underlying space which is endowed with an *ultrametric*. This opens the way for methods from ultrametric analysis in spatial information science. As a first application, we formulate a characterisation of hierarchical chain complexes in terms of a condition for 'integrating' the boundary operator on one level of detail up to a higher level. This is actually a direct consequence of the consistency rule for chain maps. A proof for that condition is defered to [1], where a formulation is provided in terms of integrals over ultrametric spaces.

2 Motivation: Segmentation and Cartographic Generalisation

Image segmentation is often hierarchic, because many classification methods are hierarchic. This leads to a tree representation of the segment hierarchy. In such a situation, topological modelling asks for a topology in each level in order to obtain an understanding of the image content. And this presumes a compatibility of the level topologies along the hierarchy. This situation is analogous to that in cartographic generalisation in which a hierarchic family of maps in various scales are produced. The main difference is that the 'segment' is now the 'region' or 'area' or some other spatial entity. Henceforth, we will use the word 'segment' as a synonym for those notions.

The hierarchical structure decomposes a segment in one level into the disjoint union of segments at the next level. The topological model consists in a boundary relation between segments in the same level. This relation can be enhanced with additional combinatorial information (like orientation or multiplicity) which in fact is again of topological nature. Two different points of view are prevalent. Namely, that of the topologist versus the topological modeller. The topologist contends himself with understanding spaces whose geometry is 'stripped off'. This abstraction from geometry reveals a coarser structure underlying the space, and which does not depend on the chosen geometry. This is the 'topologist's topological information'. The modeller wants to know more. He wants also to know the role of each piece the given space is made of in producing that kind of information. This implies a serious interpretation of the boundary relation as the set of object pairs actually being adjacent, the 'modeller's topological information'. In contrast, the topologist is more interested in the quality of the boundary relation than its exact object-object specification.

In order to obtain a topological understanding of images (or other data), the modeller will aim at deriving a topological model of the scene. From an image processing point of view, it makes sense for the model to be a hierarchical one. Further, it should

allow modifications and consistency checks. This motivates us to take the road from images to maps through hierarchy. In the following section, we review the topologist's and the modeller's points of views on topological notions at one single level: from topological datatypes to Euler-Poincaré operators. The latter are an n-dimensional generalisation of the known Euler operators from volume modelling. In fact, these are 'topologist's modifications', because they do not actually say which parts of the space are modified—they only quantify the changes in the combinatorial and topological data of a given space. Hence, they are useful for topological consistency checking. The hierarchical viewpoint in the next section necessitates to formulate everything in a relative setting $X \to Y$. This results in operators for mappings between topological datatypes. This is followed by a section on the hierarchic versions of datatypes and Euler-Poincaré operators. The hierarchy allows for defining a range of such 'topologist's topological modifications' without having to lay hands on the actual cells. Naturally, an Euler operator cannot 'exist' without an actual modification. Therefore, these are also treated in this context. The result is a metric on the space of Euler-Poincaré operators, which in fact turns out to be an ultrametric in the last section. Ultrametrics are made for measuring distances within hierarchies. After a short introduction to ultrametric spaces, we show how to 'lift' a given boundary operator in the level of highest 'resolution' in a hierachy up to other levels. This is depends on a so-called *integrability condition* which we derive. This condition describes in fact the 'compatibility' of the topologies between different levels in the hierarchy. The name 'integrabiliy' is, in fact, not chosen arbitrarily: we will give an idea on how the entries of the lifted boundary matrix can be computed by integrating on the ultrametric space.

3 Euler-Poincaré Formula and relational chain complexes

In this section we review the topological datatypes leading to chain complexes and their relational realisations as well as Euler-Poincaré operators. We adopt the viewpoint of the topologist as well as that of the topological modeller.

3.1 Graphs as low-dimensional topological datatype

The first known topological datatype is the graph encoding adjacency of regions. Through this formalisation, Euler was able to solve the Königsberg bridge problem [6]. Its vertices correspond to the regions, and every edge to the connection of two regions by a bridge. It was the abstraction from geometry to the topology encoded in the graph which allowed to find a combinatorial solution to the problem.

In cartography, not only the adjacency of regions through lines, but also the adjacency of lines through points is of importance. Thus, a map is encoded by a planar graph. Often a planar graph itself is encoded by providing each edge e with an orientation and then saying which points are 'in front of' or 'behind', and which regions are 'left' or 'right' of e. This idea stands behind the datatype DIME.

The initial topological information provided by a graph (or planar graph) is connectivity: the topological space X it encodes is connected by paths if and only if the graph is connected. Even more, the number $b_0(X)$ of connected components of X equals the number of connected components of the graph. It is clear that this number depends only on the topology of X, and not on any geometry this space may have. Even furthermore, the number b_0 does not change if X undergoes any topological transformation $X \to Y$ which has a topological inverse transformation $Y \to X$. A topological transformation is nothing but a *continuous map*, and a transformation having a continuous inverse is called *homeomorphism*. Any quantity associated with X, which does not change under any homeomorphism, is called a *topological invariant*. Hence, the connectivity number $b_0(X)$ is a topological invariant. It is named the 0-*th Betti number* of X. In the graph case, we speak of *graph isomorphism*. If there are labels on the vertices or edges, an isomorphism must respect the labels, i.e. can only identify equally labelled edges or vertices.

Topological invariants provide a first idea about the shape of a space. The connectivity number compares space X with a set of points. In fact, if the space X can be contracted to a point without loss of topological information (in the topologist's sense[3]), then X is not distinguishable from a point. This gives already for graphs an idea on how to determine the Betti number b_0: contracting all edges which are not themselves loops results in a graph which looks like a bouquet of flowers in a meadow, each connected component consisting of one single vertex with some loop-edges attached to it. The number of vertices now equals b_0. The graph contraction method above reveals another topological invariant: the number of (minimal) circuits, given by the loop-edges in the contracted graphs. This is the first Betti number b_1. In a similar manner, b_2 defines the number of cavities enclosed by a surface. For example, a planar graph can in fact be viewed as a graph on a sphere, if the 'outer face' is completed by a point at infinity. Hence, $b_2 = 1$ for planar graphs.

The topological invariants are related to the combinatorial data. Already for the contracted graphs, it can be observed that the number V of vertices, and the number E of edges are in the following relation

$$V - E = b_0 - b_1, \tag{3.1}$$

and this holds true for any graph, because general circuits have the property $V = E$. For planar graphs, this relation becomes

$$V - E + F = b_0 - b_1 + b_2 = 2, \tag{3.2}$$

where F is the number of faces. Notice that the number of 'inner' faces equals the number of minimal edge circuits. In other words, all these circuits are 'filled in' by a face and hence become contractible. This means that $b_1 = 0$. Together with $b_0 = b_2 = 1$, equation (3.2). In the intermediate case, some circuits are 'filled in' by a face, others not. In this case, b_1 can be positive.

[3] more to this in subsequent sections

The system behind this approach is the *boundary relation*. A face has a number of lines at its boundary, a line has points, and a point has empty boundary. Clearly, this homological approach can be generalised to arbitrary dimension: the cell complex. Hence, we obtain a very simple topological datatype which can handle any finite dimension. Its simplicity lies in a structure which is independent of dimension. There, the circuit generalises to a cycle of n-cells which potentially enclose an $n+1$-cell. The crucial point is to identity those cycles which are not a boundary. In the following section, we will see how this idea leads to an algebraic formalisation of the cell complex by allowing for multiplicities. As an example, take the loop: it is an edge having at its boundary one single point, because the vertex at its origin coincides with its terminal vertex. This equality means algebraically that the difference between these two vertices is zero. Through this algebraisation, the Betti numbers become dimensions of spaces which parametrise cycle classes. And computational topology has found an algorithm which allows to track cycle classes along a growth process of cell complexes in order to compute a topological barcode of such a dynamic structure.

3.2 Topological datatypes for topologists

An important notion in topology is that of a *cell*. This is a topological space X which 'looks like' an (open) ball B_r in some Euclidean space $\mathbb{R}^n$. In this case, X is an *n-cell*. The term 'looks like' is topologically defined as *homeomorphic*, meaning that there are continuous maps $X \to B_r$ and $B_r \to X$ which are inverse to each other. An intuitive description often encountered is that of X being a deformation of the ball B_r. Cells are the building blocks of *cell complexes* or *cw-complexes* which are obtained recursively by attaching an $n+1$-cell along its boundary (an *n-sphere*) to the *n-skeleton*. A good feature of a cw-complex X is that it allows an algebraic description. Namely, by saying that a *chain* (reminiscent of 'path') is a formal linear combination of cells. The idea is that of traversing along a path, one passes through a number of cells: perhaps some are encountered more than once, and others perhaps opposite to their given orientations. If all cells in a chain are of the same dimension n, one speaks of an n-chain. The space of all n-chains is denoted as

$$C_n(X) = \mathbb{Z}B_n = \bigoplus_{b \in B_n} \mathbb{Z}b,$$

and means (in various notational forms) the module of all positive and negative sums of n-cells from B_n.

Saying which n-cells are at the boundary of a given $n+1$-cell together with their multiplicities yields the *boundary operator*

$$\partial x = \alpha_1 y_1 + \ldots \alpha_m y_m,$$

for every n a linear map $\partial\colon C_{n+1}(X) \to C_n(X)$ on the whole chain module $C(X) = \mathbb{Z}B = \bigoplus_n C_n(X)$. This defines the *chain complex* $\mathcal{C}(X) = (C(X), \partial)$ associated to X.

One can derive the important property

$$\partial^2 = 0 \tag{3.3}$$

which is the algebraic equivalent of saying that the boundary of a cell itself has no boundary. Observe that the boundary operator has the same structure in every cell dimension n, where it is a linear operator ∂_n satisfying $\partial_n \partial_{n+1} = 0$, the 'local' analogue of (3.3). In other words, the chain complex looks in dimension n not much different than in dimension 2 or 1. Only that for graphs it has the special form $\partial(e) = P - Q$.

The topologist's motivation for discovering the chain complex is that it allows to compute topological invariants algebraically (at least in principle). So, our topologist has in fact become an *algebraic topologist*. For example, in a graph Γ it can be decided whether an edge is a loop or not by looking at its boundary. Namely, $\partial e = 0$ if and only if it is a loop. The same holds true for an arbitrary 1-chain $c \in C_1(\Gamma)$: it represents a cycle, if and only if $\partial c = 0$, as can be easily verified by modifying the loop through the insertion of further edges. The observation that a sphere has no topological boundary yields the generalisation to higher dimensions. Hence, a sphere is a 2-cycle. Therefore, the *cycle space* in a chain complex is the null space (kernel) of the boundary operator:

$$\{\text{cycles}\} = \ker \partial = \{x \in C(X) \mid \partial x = 0\},$$

a linear space which can be determined by methods from linear algebra. Next, observe that the topological boundary of a ball is a sphere, and this is a cycle. In other words, a boundary itself is a cycle:

$$\operatorname{im} \partial \subseteq \ker \partial, \tag{3.4}$$

where $\operatorname{im} \partial = \{\partial x \mid x \in C(X)\}$ denotes the image space of ∂ (i.e. the space of all boundaries). The number of loops in a planar graph is now determined as the number of minimal 1-cycles minus those which are boundaries of 2-cells. Since the cycle space $\ker \partial$ always has a basis consisting of minimal cycles, we infer the rule:

$$b_1(\Gamma) = \dim(\ker \partial_1) - \dim(\operatorname{im} \partial_2),$$

and the general Betti numbers are defined in this way.

Definition 1. *Let C be a chain complex. Then*

$$b_n(C) = \dim(\ker \partial_n) - \dim(\operatorname{im} \partial_{n+1})$$

is called the n-th Betti number of C. *If $C = C(X)$ for some cellular space, then $b_n(X) := b_n(C(X))$ is the n-th* Betti number of X.

Observe that the Betti numbers formalise the intution of spaces having a certain number of n-dimensional 'holes'. And b_0 is the number of connected components of X. For a large class of spaces X the Betti numbers obtained from different kinds of

associated chain complexes coincide. The observation (3.4) is equivalent to the condition (3.3) imposed on the chain complex. But there is more to it: the usual construction in linear algebra for a situation as in (3.4) is to define the quotient module

$$H_n(X, \mathbb{Z}) = \ker \partial_n / \operatorname{im} \partial_{n+1},$$

obtained by identifying those n-cycles whose differnce is a boundary. Then, $b_n(X) = \dim H_n(X, \mathbb{Z})$. Also, sometimes it is necessary to consider individual cycles upto some boundary.

One way of comparing an object X with another object Y is by constructing a mapping $f\colon X \to Y$ and to study its properties. In the case of topological spaces, it makes sense to require that f be a continuous map, because then differences between the topological properties of X and Y can be found. If X and Y have a cellular structure, then f induces a linear map $C(X) \to C(Y)$ by saying that the image of a cell c in X is some linear combination of the cells of Y containing $f(c)$. If f is a continuous mapping between cw-complexes which is *cellular*, i.e. maps n-cells of X into the n-skeleton of Y, then there is an induced linear map $f_*\colon C_n(X) \to C_n(Y)$ such that cycles map to cycles and boundaries to boundaries. In other words, there is a linear map

$$f_*\colon H_n(X, \mathbb{Z}) \to H_n(Y, \mathbb{Z}).$$

The corresponding map on chain complexes is also called *cellular*, or a *chain map*, in this case. One application of this fact is that if there is a surjective cellular map $f\colon X \to Y$, then the Betti numbers of Y cannot be larger than those of X.

In the following section, we will have chain maps $C(X) \to C(Y)$ obtained from 'segmentation mappings' $Y \to X$ which for a given segmentation $A_1 \sqcup \cdots \sqcup A_n = A$ take the segments A_i to the generalised cell A. The linear map on chain complexes $C(X) \to C(Y)$ then maps A to a linear combination of the A_i, which yields a cellular map.

3.3 Topological datatypes for n-dimensional modelling

The chain complex is the topologist's ideal data type. From it the relevant topological information on X can be read of: the Betti numbers, or, if a closer look is necessary, the homology or cohomology[4] of a given space X. For the modeller, however, the chain complex is not sufficient, because he wants to lay his hands on its individual constituents (called *cells* by abuse of language). The topological invariants are not of primary interest, of concern is the precise boundary relation. The notion of *relational chain complex* is a datatype which can do both, boundary relation (the *modeller's topology*) and homology (the *algebraic topologist's topology*). The idea is, in fact, simple: model the boundary relation as a graph, add the coefficients of the boundary operator as weights, and define a multiplication such that the important rule $\partial^2 = 0$ holds

[4]This is a dual notion to homology, and will not be used in the remainder of this article.

true for these new modeller's boundary operators. This idea has a simple realisation with partial matrices, for which a natural multiplication can be defined. The following definition is an immediate consequence of this idea:

Definition 2. *A* relational chain complex *is a pair $\mathfrak{C} = (B, D)$, where B is a finite set partitioned into subsets $B_0, \ldots, B_n$, and D is the graph of a partial matrix $\partial :\subseteq B \times B \to \mathbb{Z}$ satisfying:*

1. *∂ is defined on $x \in B_j \times B_i$ only if $j < i$.*
2. *If $x \in B_j \times B_i$ with $j < i - 1$, then $\partial(x)$ is either 0 or undefined.*
3. *$\partial^2 = 0$ is a partial zero matrix (i.e. zero where defined).*

∂ is called the relational boundary *of $\mathfrak{C}$.*

Condition 1. means that the boundary of an i-cell can contain only cells of smaller dimension. 2. means that only cells of dimension $i - 1$ have a non-zero count at the boundary of an i-cell. The condition 3. is a natural generalisation of the topologist's condition on the boundary operator.

With the same idea as above, the definition of a relational complex morphism is almost straightforward:

Definition 3. *Let $\mathfrak{C} = (B, D)$ and $\mathfrak{C}' = (B', D')$ be relational chain complexes. A* morphism $M \colon \mathfrak{C} \to \mathfrak{C}'$ of relational chain complexes, *sometimes also called* relational chain map, *is the graph of a partial matrix $m :\subseteq B' \times B \to \mathbb{Z}$ satisfying the conditions:*

1. *If $i \neq j$ and $x \in B'_j \times B_i$, then $m(x)$ is either zero or undefined.*
2. *$m \cdot \partial - \partial' \cdot m = 0$ (partial zero matrix!).*

If instead of 2. the stronger identity

2′. $m \cdot \partial = \partial' \cdot M$

is valid, then M is called strict.

The composition $M' \circ M$ of M with $M' \colon \mathfrak{C}' \to \mathfrak{C}''$ is defined as the graph of the partial matrix $m' \cdot m :\subseteq B'' \times B \to \mathbb{Z}$.

The morphisms defined in [3] are in fact the strict morphisms in the sense of the above Definition. Condition 2. is motivated by the following example:

Example 1. *A morphism of relational complexes which is not strict is shown in Figure 1. The corresponding relational boundary operators (left: δ, right: ∂) and the morphism M are given by their partial matrices in Table 1. That this morphism is not strict is shown in Table 2.*

It is now clear that a relational chain complex $\mathfrak{C}$ (good for modellers) has an associated chain complex (the topologist's favourite). Namely, by replacing the relational boundary of $\mathfrak{C}$ by the matrix obtained from ∂ by 'filling up with zeros', i.e. the non-defined entries are all replaced by 0. This is clearly a valid boundary operator, and

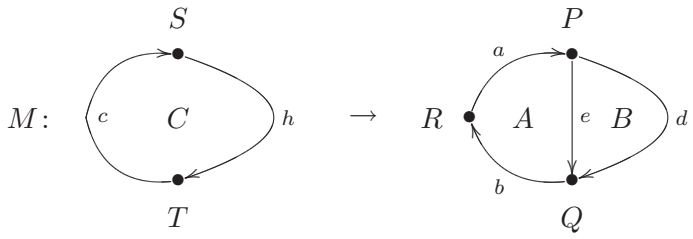

Figure 1: A non-strict morphism of relational chain complexes.

δ	C	c	h
c	1		
h	1		
S		1	-1
T		-1	1

∂	A	B	a	b	d	e
a	1					
b	1					
d		1				
e	1	-1				
P			1		-1	-1
Q				-1	1	1
R			-1	1		

M	C	c	h	S	T
A	1				
B	1				
a		1			
b		1			
d			1		
e	0				
P				1	
Q					1
R		0			

Table 1: The boundary operators and morphism of Figure 1.

the resulting chain complex $\mathcal{C}$ has 'the same topology' (for algebraic topologists!) as the relational chain complex $\mathfrak{C}$. This statement seems meaningless, as homology of relational chain complexes is not yet defined, but it is natural to define $H_n(\mathfrak{C}, \mathbb{Z})$ as $H_n(\mathcal{C}, \mathbb{Z})$. However, it is also possible to define the kernel and image of a partial matrix, and then make the trivial observation that there is a surjective map

$$t\colon H_n(\mathfrak{C}, \mathbb{Z}) \to H_n(\mathcal{C}, \mathbb{Z})$$

given by replacing 'undefined' by zero (t stands here for 'trivial'). Notice, that this makes use of the kernel being the pre-image of any partial zero vector. If relational homology is defined with the *strict kernel* (all vectors mapping to the full zero vector), then the map t is an isomorphism.

A relational chain complex $\mathfrak{C} = (B, D)$ has a natural topology on B coming from the relational boundary ∂. Namely, the property

$$'\partial(b, c) \quad \text{is defined}'$$

defines a relation Δ on B. Then the smallest topology on B for which the stars

$$S(x) = \{x_0 \Delta x_1 \dots x_{n-1} \Delta x_n \mid x_0 = x,\ x_1, \dots, x_n \in B\}$$

∂M	C	c	h
a	1		
b	1		
d	1		
e	0		
P	0	1	-1
Q	0	-1	1
R		0	

$M\delta$	C	c	h
a	1		
b	1		
d	1		
e			
P		1	-1
Q		-1	1
R	0		

Table 2: The difference $\partial M - M\delta$ is a partial zero matrix, but $\partial M \neq M\delta$.

are open is the *cell topology* on B.

Example 2. *Observe that in Example 1 there is a continuous map between the underlying topological spaces in the reversed direction given by*

$$A, b, e \mapsto C,\ a, b, R \mapsto c,\ d \mapsto h,\ P \mapsto S,\ Q \mapsto T.$$

3.4 Euler-Poincaré operators

Euler observed that for the surface of a convex polyhedron the identity

$$V - E + F = 2$$

holds true, where V, E, F are the numbers of vertices, edges, and faces, respectively. This is the same as (3.2) which generalises equation (3.1) for graphs to the case of planar graphs. The relation between the combinatorial data and topological invariants generalises to arbitrary finite N-dimensional cw-complexes:

$$\sum_{i=0}^{N}(-1)^i V_i = \sum_{i=0}^{N}(-1)^i b_i(X), \tag{3.5}$$

where V_i is the number of i-cells of X. This result is in fact true for finite chain complexes, and hence for relational chain complexes.

The relation (3.5) gives us a consistency rule for topological modifications. By this we mean that a given complex X can be modified by inserting and deleting cells and redefining a boundary operator only in such a way that the resulting object is again a valid complex. Any such modification has an immediate (non-unique) effect on the Betti numbers by making or deleting holes or connected components. Equation (3.5) then makes sure that the balance is held before and after the modification. A modification which respects that balance equation is called a *Euler-Poincaré operator*. It is represented by a $(2N + 2)$-tuple of integers $(V_0, \ldots, V_N, b_0, \ldots, b_N)$ satisfying the condition (3.5). Euler-Poincaré operators are used in volume modelling ($N = 3$) under the name *Euler-Operators*. Hence, their generalisation to N dimensions is very natural.

A general Euler-Poincaré operator is given as an integer solution of the linear equation (3.5). In other words, it is a $\mathbb{Z}$-linear combination of $2N+1$ basis solutions which we call *basic Euler-Poincaré operators*:

$$K_i = X_i - (-1)^i X_0, \qquad i = 1, \ldots, n \tag{3.6}$$

$$L_j = Y_j + (-1)^j X_0, \qquad j = 0, \ldots, n \tag{3.7}$$

where X_i and Y_i have the meaning:

$$\begin{aligned} X_i &: \text{'make' an } i\text{-dimensional 'cell'} \\ Y_j &: \text{'make' a } j\text{-dimensional 'hole'} \\ -X_i &: \text{'kill' an } i\text{-dimensional 'cell'} \\ -Y_j &: \text{'kill' a } j\text{-dimensional 'hole'} \end{aligned}$$

However, the meaning of Y_0 is of course 'join' two 'parts'.

Some examples of basic Euler-Poincaré operators are given in the following example:

Example 3. *Figure 2 illustrates some examples of* $L_0 = Y_0 + X_0$ (make vertex, make component: mvmc)*,* $L_1 = Y_1 - X_0$ (kill vertex, make loop: kvml)*,* $L_2 = Y_2 + X_0$ (make vertex, make shell: mvms) *as well as* $K_1 = X_1 + X_0$ (make vertex, make edge: mvme) *and* $K_2 = X_2 - X_0$ (kill vertex, make area: kvma)*.* L_2 *takes the Betti numbers* $b_0 = 1$*,* $b_1 = b_2 = 0$ *to* $b_0 = b_2 = 1$*,* $b_1 = 0$*. And* K_2 *leaves* $b_0 = b_1 = 1$*,* $b_2 = 0$ *unchanged.*

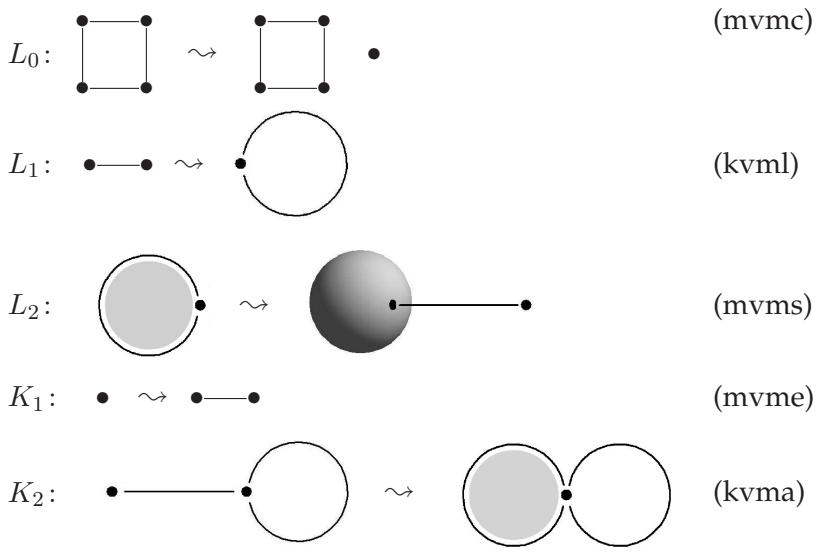

Figure 2: Examples of Euler-Poincaré operators.

Remark 1. *In the case that only oriented manifolds without boundaries are allowed, it is known that the Betti numbers are palindromic. Hence, this case has another linear constraint:*

$$b_i - b_{N-i} = 0,$$

where N is the dimension of the manifold.

One useful property of Betti numbers is their behaviour in the context of generalisation. More concretely, assume there are complexes given with cellular maps

$$A \to C \to B \tag{3.8}$$

such that $f\colon A \to C$ is injective, $g\colon C \to B$ surjective, and $\operatorname{im} f = \ker g$. Then the sequence (3.8) is called a *short exact sequence*. It is then a fact that the so-called *Euler-Poincaré characteristic*

$$\chi = \sum_{i=1}^{N} (-1)^i b_i$$

adds up:

$$\chi(C) = \chi(A) + \chi(B) \tag{3.9}$$

In the generalisation situation we consider, A is a generalisation of C, and B measures the 'difference' (in form of the isomorphy of B to the quotient C/A). Sometimes, the converse is also useful: model the detailed version of a complex B as some C with a surjective morphism to B. Then A, as the kernel of g, measures the 'loss of topological information'. Together with the obvious balance equation for cells:

$$V_i(C) = V_i(A) + V_i(B),$$

equation (3.9) yields a consistency rule for *relative Euler-Poincaré operators* given as a linear subspace in $\mathbb{Z}^{3(2N+1)}$ of dimension $6N + 2$. A basis, called *basic relative Euler-Poincaré operators*, is given by:

$$K_i^A, K_i^B, K_i^C;\ L_0^C + L_j^A,\ L_0^C + L_j^B,\ L_0^C + L_k^C \quad (i, k = 1, \ldots, n;\ j = 0, \ldots, n), \tag{3.10}$$

where the superscript denotes which of the three complexes A, B, C the operator acts on.

There exists a case, when the Betti numbers themselves add up:

$$b_i(C) = b_i(A) + b_i(B) \tag{3.11}$$

for all $i = 0, \ldots, N$. Namely, when the mappings in the short exact sequence (3.8) have reverse mappings $f'\colon C \to A$, $g'\colon C \to B$:

$$A \leftarrow C \leftarrow B$$

such that f' is a left inverse to f (i.e. $f'f = \mathrm{id}_C$) and g' a right inverse of g (i.e. $gg' = \mathrm{id}_C$). Then the short exact sequence (3.8) is said to be *split*. We obtain $5N + 2$ basic Euler-Poincaré operators for the split case:

$$K_i^A, K_i^B, K_i^C;\ L_j^A + L_j^C,\ L_j^B + L_j^C \quad (i = 1, \ldots, N;\ j = 0, \ldots, N). \tag{3.12}$$

Example 4. *A non-split example is given by the embedding of a circle (1-complex) into a disk (2-complex) as in Figure 3. The boundary operators map everything to zero (times the boundary object), except in the middle, where the boundary of the area is the line (with multiplicity one). The left arrow takes the vertex to the vertex, and the line to the line. The right arrow*

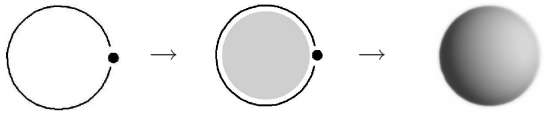

Figure 3: A non-split exact sequence.

makes the vertex and line dissapear, i.e. maps these to zero. This can be viewed as an algebraic version of contraction. The latter has no right inverse, as that would have to take the boundary of the right area (zero) to the boundary of the middle area (1 times line). This is not possible for linear maps.

Indeed, we observe that the Betti numbers b_1 and b_2 do not add up: $b_0 = 1, 1, 0$, and b_1 is $1, 0, 0$ from left to right, and b_2 takes the values $0, 0, 1$. However, the Euler-Poincaré characteristic χ is $0, 1, 1$ and adds up indeed.

Notice that the image on the right of the sequence in Figure 3 depicts an algebraic version of a 'sphere' without any further decomposition into 'cellular parts'. Therefore, the homology of that complex is different from the homology of a genuine topological sphere. The latter has $b_0 = 1$, but our 'sphere' here does not have any 'connected component' because $b_0 = 0$.

4 Spaces of Hierarchical Relational Complexes

In this section we consider complexes of fixed dimension N.

An Euler-Poincaré operator as introduced in Section 3.4 can be viewed as a topologist's modification on chain complexes. However, from the modeller's point of view, this 'operator' provides merely a statistic of some modeller's modification. In what follows, we construct hierarchical Euler-Poincaré operators which are very much like actual modifications on complexes.

Picking up the idea of generalisation or multi-resolution, we start with a hierarchical tree of 'segments' and wish to construct topologies for each level. The crucial point is that the topologies at different levels should be compatible in some way in order to be able to compute the topological differences between levels. Further, modifications

must be made in a consistent manner: a modification in one level of detail must correspond to some (possibly trivial) modification in all other levels. For this aim, the preceding section provides some answers.

Assume that each 'segment' in level ν is the disjoint union of 'segments' in level $\nu+1$. This yields a tree of 'segment' hierarchies, if we assume that the top level $\nu = 0$ is the 'generic segment': the disjoint union of all segments of level $\nu = 1$ (or any fixed level ν). Next, we assume that the set X_ν of all 'segments' at level ν bears a topology in the form of a relational complex. In other words, each 'segment' has a dimension and is in a boundary relation with some other 'segments'. The task is to find a suitable way of defining the relational boundary operator, compatible with the hierarchical structure on all 'segments'.

First of all, observe that there is a mapping $\pi_{\nu,\nu+1}\colon X_{\nu+1} \to X_\nu$ which takes all children 'segments' to their common parent 'segment'. The first condition we ask for is that $\pi_{\nu,\nu+1}$ be continuous. Another condition is that if $x \in X_\nu$ is of dimension n, then $\pi^{-1}_{\nu,\nu+1}(x)$ be open in the n-skeleton of $X_{\nu+1}$. This requirement makes sure that if x is viewed as an open subset of X_ν, then its representation as the union of its children is also open in $X_{\nu+1}$. Under these conditions, we can define a linear map between the chain spaces:

$$\psi_n^{\nu,\nu+1}\colon C_n(X_\nu) \to C_n(X_{\nu+1}), \quad X_\nu \ni x \mapsto \sum_{y\in\pi^{-1}_{\nu,\nu+1}(x)} \alpha_y y, \tag{4.1}$$

which takes each n-cell of X_ν to a linear combination of its n-dimensional parts in $X_{\nu+1}$ such that:

$$\alpha_y \neq 0 \Rightarrow \dim(y) = \dim(x),$$

just like in Example 1. This defines maps

$$\psi^{\nu,\nu+1}\colon C(X_\nu) \to C(X_{\nu+1}), \quad \psi_n^\nu\colon C_n(X_\nu) \to C_n(X), \quad \psi^\nu\colon C(X_\nu) \to C(X),$$

where X is the 'segmentation' at the highest level of detail: $X = X_D$, where D such that $X_\nu = X_D$ for $\nu \geq D$. The reverse map on the underlying topological spaces generalises Example 2. Together with the trivial map $C(X_0) \to C(X_1)$ which takes the 'root' to zero, we obtain a sequence $\psi = (\psi^\nu)$ of cellular morphisms of chain complexes. We call this sequence a *hierarchical chain complex*.

From the modeller's point of view, it is now a simple task to leave out the 'spurious' zeros in order to obtain relational chain complexes $\mathfrak{C}(X_\nu)$ and inter-level morphisms. Namely, the relational boundary operators are defined only on pairs of adjacent 'segments' (of the right dimensions). And the morphism $\Psi^{\nu,\nu+1}\colon \mathfrak{C}(X_\nu) \to \mathfrak{C}(X_{\nu+1})$ between levels is defined precisely where the α_y in (4.1) are defined, and there takes the value α_y. We call the sequence $\Psi = (\Psi^\nu)$ a *hierarchical relational chain complex*.

Of interest for modellers is the following:

Lemma 1. *Let Ψ be a hierarchical relational chain complex. Assume that the chain map ψ^ν from level ν is injective. Then the boundary operator at level ν is uniquely determined by the morphism Ψ_ν upto possible 'spurious' zeros.*

Proof. By cellularity, it holds true that

$$\psi_\nu \partial_\nu = \partial \psi_\nu, \tag{4.2}$$

where $\psi_\nu \colon C(X_\nu) \to C(X)$ is the associated morphism of cell complexes and ∂_ν, ∂ are the corresponding boundary operators, all obtained by 'filling up with zeros' (cf. Section 3.3). The map ψ_ν is injective, hence there is at most one solution matrix ∂_ν, if ψ_ν and ∂ are known. In fact, (4.2) is a linear equation for each column of ∂_ν. □

The point is that the hierarchical 'segmentation' needs only to say in which way parent 'segements' are represented as weighted sums of children 'segments'. Then a boundary operator on the highest level of detail yields in a unique way the boundary operators at each other level. Thereafter, unnecessary zeros can be eliminated in order to obtain relational boundary operators.

Now observe that a hierarchical relational chain complex Ψ is given by the following data:

1. a rooted labelled tree $\mathcal{T}$
2. a labelled directed acyclic graph ∂_ν on each level ν

fulfilling the requirements above. For 1. the root is X_0, vertices are labelled by dimension, and edges by the morphisms' partial matrix entries. The relational boundary operator is as in 2.

We remark that the injectivity of ψ^ν can be checked right on the labelled tree. This is the case, if and only if each vertex has at least one edge downwards with non-zero label.

An *isomorphism of hierarchical relational complexes* is given by an isomorphism of the rooted labelled tree together with isomorphisms of labelled directed acyclic graphs on each level. This allows to consider the space of all such things upto isomorphy:

Definition 4. *The* space $\mathcal{M}_n$ of all hierarchical relational complexes on n objects *is the set of all isomorphy classes of hierarchical relational complexes whose rooted labelled tree has precisely n terminal vertices.*

In $\mathcal{M}_n$ the levels ν can take all values from $\mathbb{N}$. As there are only finitely many terminals in the rooted tree, there is for each $x \in \mathcal{M}_n$ a highest level of detail D as defined above. Notice that the number D can be arbitrarily large, because trivial inter-level maps are allowed.

Remark 2. *The fixed number n provides no practical restriction, as it can be taken very large. The subset of actually modified objects is then contained in the set of n possible objects.*

A *modification* is nothing but a self-map $F \colon \mathcal{M}_n \to \mathcal{M}_n$. It is clear that a modification F defines for each given $x \in \mathcal{M}_n$ in each level ν an Euler-Poincaré operator. This can be read off the pair $(x_\nu, (Fx)_\nu)$, where the subscript ν denotes 'level ν'.

Assume now that some set $\Sigma \subseteq \mathcal{M}_n$ of hierarchical relational chain complexes is given, and a modification F. If in level ν this Euler-Poincaré operator is the same for all $x \in \Sigma$, then we say that the restriction $F|_\Sigma$ of F to Σ has an *associated Euler-Poincaré operator at level* ν.

Definition 5. *A* hierarchical Euler-Poincaré operator on Σ *is a sequence* $\Phi = (\Phi_\nu)$ *of Euler-Poincaré operators at level* ν *which are associated to the same restricted modification* $F|_\Sigma$. *A modification* F *which defines a hierarchical Euler-Poincaré operator on* Σ *is called* Σ-rigid.

Example 5. *Any modification* F *is* $\{x\}$*-rigid for* $x \in \mathcal{M}_n$. *Quite often, one starts with a single* $x \in \mathcal{M}_n$, *and creates some new candidate complex* Fx, *and uses Euler-Poincaré operators to check if* Fx *is a valid complex. However, our concept of* Σ*-rigdidity allows for a more systematic way of altering hierarchical complexes.*

The set of all Σ-rigid modifications $F\colon \mathcal{M}_n \to \mathcal{M}_n$ has a metric:

$$d_\Sigma(F,G) = 2^{-\text{maximal level } \nu \text{ at which } F|_\Sigma \text{ and } G|_\Sigma \text{ coincide}},$$

and this allows to define the Σ-*range*

$$\rho_\Sigma(F) := d_\Sigma(F, \mathrm{id}),$$

where id is the identity map (trivial modification). The range indicates the level up to which $F|_\Sigma$ is trivial.

The hierarchical Euler-Poincaré operators on Σ form a $\mathbb{Z}$-module $\mathscr{E}(\Sigma)$ by componentwise addition, and with zero element $\mathbf{0}$, the zero vector. There is a norm, called *range*, defined as

$$\|\Phi\|_\Sigma := 2^{-\text{maximal level } \nu \text{ at which } \Phi \text{ coincides with } \mathbf{0} \text{ on } \Sigma}, \tag{4.3}$$

and a translation invariant distance:

$$d_\Sigma(\Phi, \Phi') := \left\|\Phi - \Phi'\right\|_\Sigma . \tag{4.4}$$

The range satisfies the following properties:

$$\|\Phi\|_\Sigma = 0 \Rightarrow \Phi = \mathbf{0} \tag{4.5}$$

$$\|m \cdot \Phi\|_\Sigma = \|\Phi\|_\Sigma \quad (m \in \mathbb{Z}) \tag{4.6}$$

$$\left\|\Phi + \Phi'\right\|_\Sigma \leq \max\left\{\|\Phi\|_\Sigma, \left\|\Phi'\right\|_\Sigma\right\} \tag{4.7}$$

In other words, $\mathscr{E}(\Sigma)$ is a normed $\mathbb{Z}$-module, where $\mathbb{Z}$ is endowed with the trivial norm.

The effect of a trivial Euler-Poincaré operator is to leave fixed both, the topological and combinatorial data. The following example illustrates this:

Figure 4: The effect of some trivial Euler-Poincaré operator on some tree.

Example 6. *A trivial Euler-Poincaré operator on graphs is given by 'slides' as in Figure 4.*

We remark that not only the intra-level modification, but also the inter-level morphisms can be considered as coming from an Euler-Poincaré operator. In the latter case, of course, the number of objects increases. However, there is another balance equation which we write down for $\Sigma = \{x\}$:

$$\Phi_\nu + \Upsilon_\nu(Fx) = \Upsilon_\nu(x) + \Phi_D, \tag{4.8}$$

where F is the modification, Φ the associated hierarchical Euler-Poincaré operator for $\{x\}$, Φ_μ the Euler-Poincaré operator at level μ,

$$D = \max\{\text{highest level of detail of } x, \text{highest level of detail of } Fx\},$$

and $\Upsilon_\nu(y)$ the Euler-Poincaré operator associated to the morphism defined by y from level ν to the highest level of detail of y.

5 Final Observation: Hierarchies are Ultrametric

In the previous section, we saw that a hierarchical relational chain complex Ψ has an underlying tree $\mathcal{T}$ with root X_0, whose vertices correspond to the 'segments'. The 'segments' corresponding to the terminal vertices could be considered as 'atoms'. Then any 'segment' at some vertex v decomposes into its atoms at the bottom of the branch downwards from v.

5.1 Ultrametric spaces

Any rooted tree can be viewed as a so-called *ultrametric* space in the following way. Any two 'atoms' x, y at the end of $\mathcal{T}$ have a vertex $v(x, y)$ from which the path from the top down to $\{x, y\}$ branches off. The level $\nu(x, y)$ determines the distance between x and y:

$$d_\Psi(x, y) := 2^{-\nu(x,y)},$$

which together with $\nu(x, x) := \infty$ is a non-negative function which satisfies the axioms for distance:

$$d_\Psi(x, y) = 0 \Rightarrow x = y \tag{5.1}$$

$$d_\Psi(x, y) = d_\Psi(y, x) \tag{5.2}$$

$$d_\Psi(x, y) \leq \max\{d_\Psi(x, z), d_\Psi(z, y)\} \tag{5.3}$$

(5.3) is called the *strong triangle inequality*, because it is stronger than the usual triangle inequality. Any distance satisfying the strong triangle inequality is called an *ultrametric*. In other words, hierarchical relational chain complexes are ultrametric spaces.

The hierarchical structure defined by $\mathcal{T}$ allows to extend d_Ψ to arbitrary vertices:

$$d_\Psi(v_1, v_2) := d_\Psi(x_1, x_2),$$

where x_i is the terminal vertex of any path from root through v_i. Notice that this extended definition of d_Ψ does not depend on the choice of x_i.

The ultrametric distance d_Ψ is a measure for the relevance of a given set S of 'atoms' or 'segments'. Namely, the larger the *diameter* of S (defined as the maximal distance between two elements of S) the more levels of detail are affected. This applies, for example, in cartographic situations when some buildings or roads in some urban situation are altered.

In Section 4, we have already encountered an ultrametric space. Namely, the range (4.3) of hierarchical Euler-Poincaré operators was used to define an ultrametric (4.4) on the space $\mathscr{E}(\Sigma)$ of hierarchical Euler-Poincaré operators on Σ. The range itself is an *ultrametric norm* on $\mathscr{E}(\Sigma)$.

5.2 ψ-integrability of boundary operators

The labelling of the tree $\mathcal{T}$ of the previous subsection can be interpreted as a kind of integration rule when some 'segment' x is decomposed into its 'atoms' $x_1, \ldots, x_r$. For instance, there is the dimension rule from vertex labels

$$\dim(x) = \max\{\dim(x_1), \ldots, \dim(x_r)\} \tag{5.4}$$

which can be inferred from the definition of Ψ. Equation (4.1) is another such integration rule coming from the edge labels. A consequence of these two integration rules is the construction of the boundary operator at level ν from a given boundary operator ∂ on the 'atomic' level. From this point of view, Lemma 1 says that an 'atomic' boundary operator has a unique lift to any level ν under the unspoken condition that a lift exists.

In order to determine conditions under which ∂ has a lift, assume that the hierarchical map ψ at level ν: $\psi_\nu \colon C(X_\nu) \to C(X)$ is given as a linear map, i.e. we assume that the boundary operator ∂_ν of $C(X_\nu)$ is unknown if existent, and define:

Definition 6. *The boundary operator ∂ is called* integrable *with respect to the linear map ψ_ν (or for short: ψ_ν-*integrable*), if there exists a boundary operator ∂_ν on $C(X_\nu)$ such that ψ_ν is a cellular map of chain complexes. If the boundary operator ∂ is ψ_ν-integrable for all ν, then it is called ψ-*integrable.

The question is now, which conditions the linear map ψ must fulfill in order for ∂ to be ψ-integrable. Here, the map $\pi_\nu \colon X_\nu \to X$ which takes a 'segment' to its 'atoms' plays an important role. Recall that, according to (4.1), ψ_ν takes an n-cell $x \in X_\nu$ to

a linear combination of n-cells in X contained in $\pi_\nu^{-1}(x)$. In general, the fibre $\pi_\nu^{-1}(x)$ contains also cells of dimension lower than $n = \dim(x)$. These are, by assumption, contained in the interior of x. Assume that c is such an interior cell, and that $\dim(c) = n-1$. Being in the interior of x means that c does not appear in the boundary $\partial\psi_\nu(x)$ in its representation as a linear combination of $n-1$-cells. In other words, the coefficient α_c in (4.1), applied to $\psi_\nu \colon C_{n-1}(X_\nu) \to C_{n-1}(X)$, vanishes, if ∂ is ψ_ν-integrable.

For a more systematic approach, we introduce some notation. The unknown boundary coefficients of x under ∂_ν will be written as $(x : c)_\nu$, where c runs through the cells of X_ν:

$$\partial_\nu\, x = \sum_c (x : c)_\nu\, c.$$

The coefficients of a chain $d \in C_n(X)$ are denoted as $\langle d \mid y\rangle$, where y runs through the cells of X. Hence, we have the expression

$$\psi_\nu(x) = \sum_y \langle \psi_\nu(x) \mid y\rangle\, y = \sum_{y \in \pi_\nu^{-1}(x)} \langle \psi_\nu(x) \mid y\rangle\, y.$$

And for a cell y of X, the expression $(y : b)$ denotes the coefficient of the given boundary operator ∂, which yields:

$$\partial y = \sum_b (y : b)\, b.$$

From the chain map definition

$$\partial\psi_\nu(x) = \psi_\nu\partial_\nu(x) \tag{5.5}$$

we obtain the following result, whose proof is deferred to [1]:

Theorem 1. *The boundary operator ∂ is ψ_ν-integrable, if and only if for all $x, c \in X_\nu$ the equations*

$$(x : c)_\nu \langle \psi_\nu(c) \mid b\rangle = \sum_{y \in \pi_\nu^{-1}(x)} \langle \psi_\nu(x) \mid y\rangle (y : b) \tag{5.6}$$

have a common solution $(x : c)_\nu$ for all b having the same image $\pi_\nu(b) = c$.

Notice that if ψ_ν is injective, then ∂_ν is a valid boundary operator. This follows from the observation:

$$\psi_\nu\partial_\nu^2 = \partial\psi_\nu\partial_\nu = \partial^2\psi_\nu = 0,$$

where the cellularity property (5.5) has been used twice. By injectivity of ψ_ν it follows now that $\partial_\nu^2 = 0$, and ∂_ν is indeed a boundary operator.

Example 7. *Consider the situation of Figure 5 with the linear map M from Table 1 and the continuous map π as defined in Example 2, with the difference that the in M one column is only partially known:*

$$M(C) = \alpha A + \beta B.$$

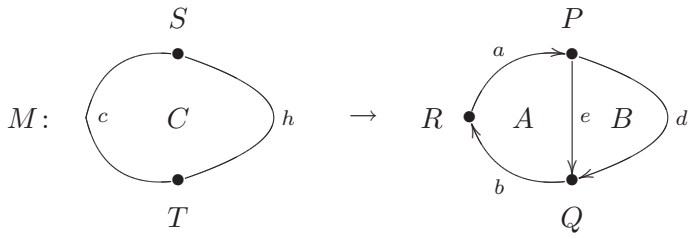

Figure 5: A linear map between relational chain complexes with unknown boundary operator δ (mind the missing arrow heads on the left).

In this example, the boundary operator δ is unknown, as indicated by the missing arrows heads in the left of Figure 5. From equation (5.6), one computes in the case $\alpha = \beta = 1$ the boundary operator δ as in Table 1, except that all blanks are filled with zeros.

We now want to find out why in Example 1 we did choose $M(C) = A + B$ in the first place, and not e.g. $M(C) = A - B$.

Answer. *Equation (5.6) translates for the unknown boundary coefficient $(C : C)$ to the equalities:*

$$\begin{aligned}(C:C)\langle M(C) \mid A\rangle &= \langle M(C) \mid A\rangle(A:A) + \langle M(C) \mid B\rangle(B:A)\\ (C:C)\langle M(C) \mid B\rangle &= \langle M(C) \mid A\rangle(A:B) + \langle M(C) \mid B\rangle(B:B)\\ (C:C)\langle M(C) \mid e\rangle &= \langle M(C) \mid A\rangle(A:e) + \langle M(C) \mid B\rangle(B:e)\end{aligned}$$

With the data from Table 1 (under our restrictions) this yields the equivalent system:

$$\begin{aligned}(C:C)\cdot\alpha &= 0\\ (C:C)\cdot\beta &= 0\\ (C:C)\cdot 0 &= \alpha - \beta\end{aligned} \tag{5.7}$$

In other words: $\alpha = \beta$, and if that is non-zero, then $(C : C) = 0$. Equation (5.7) is responsible for our choice $\alpha = \beta = 1$ in Example 1. For a choice $\alpha \neq \beta$, equation (5.7) becomes a contradiction, and ∂ is not M-integrable. The possibility $\alpha = \beta = 0$ yields a degree of freedom for the value of $(C : C)$. In that case, $(C : C) = 0$ has to be imposed for the sole reason of forcing δ to be a boundary operator.

Assume that ψ belongs to a hierarchical relational complex Ψ. Define a relation Δ_ν on X_ν by declaring $\pi_\nu(x)\Delta_\nu\pi_\nu(y)$, if $x\Delta y$. Here, Δ is the relation underlying the relational boundary operator ∂ as introduced in Section 3.3. We denote the relational boundary operators of X and X_ν coming from Ψ by δ and δ_ν. Equation (5.6) can be rewritten as

$$\delta_\nu(c,x)\langle\Psi_\nu(c) \mid b\rangle - \sum_{y\in\pi^{-1}(x)} \langle\Psi_\nu(x) \mid y\rangle\delta(b,y) = 0 \text{ or undefined,}$$

with $b \in \pi_\nu^{-1}(c)$, and the expressions in $\langle\ \rangle$ are the entries of the partial matrix belonging to the relational morphism Ψ_ν. The relation Δ_ν is 'good' in that its domain contains all pairs for which δ_ν is non-zero, provided $\langle \psi_\nu(c) \mid b \rangle \neq 0$. Namely, if $c\Delta x$ is not defined, then the right hand side of (5.6) is zero. It follows then from (5.6) that $(x : c) = 0$.

The following corollary is an algebraic way of saying that any $n-1$-dimensional part of an n-cell x lies in the interior of x, and is proven in [1].

Corollary 1. *Assume that there is some $e \in \pi_\nu^{-1}(x)$ with $\dim(e) = n-1$ and $x \in X_\nu$ with $\dim(x) = n$. If ∂ is ψ_ν-integrable, then $\langle \partial\psi_\nu(x) \mid e \rangle = 0$.*

We further remark ψ-integrability is equivalent to integrability defined by the ultrametric inter-level 'integration rule' $\psi^{\nu,\nu+1} \colon C(X_\nu) \to C(X_{\nu+1})$ for all levels ν. The 'integration' itself is given by equation (5.6) which could in a suggestive fashion also be written as

$$(x : c) = \int\limits_{\pi_\nu^{-1}(x)} (y : b)\, d_b y$$

with the 'measure' $d_b y = \frac{\langle \psi_\nu(x) | y \rangle}{\langle \psi_\nu(c) | b \rangle}$ for $b \in \pi_\nu^{-1}(c)$ such that the expression $\langle \psi_\nu(c) \mid b \rangle$ is non-zero. This idea is made rigourous in [1].

The consequence of ψ-integrability is that only a ψ-integrable boundary operator at the highest level of detail is needed. Assume that a hierarchical complex is modified by an instance of an Euler-Poincaré operator at the highest level, and then along the new hierarchy a hierarchical linear map ψ' is constructed. If the new boundary operator ∂' is ψ'-integrable, then it determines the boundary operators at all levels together with the corresponding hierarchical Euler-Poincaré operator.

Remark 3. *A hierarchical chain complex ψ is an instance of a* persistence complex, *as defined in [10]. Such a complex allows the computation of a* persistence barcode *which represents the 'lifetimes' of basis elements of the homology groups $H_n(C(X_\nu))$ across the values of ν in a multiset of intervals. In other words, a ψ-integrable boundary operator ∂ allows for persistence barcodes. The issue in generalisation from segmentation thus lies in the construction of pairs ψ, ∂ such that ∂ is ψ-integrable. Hence, a* relational persistence complex *is defined to be a sequence of relational chain maps $\mathfrak{C}^\nu \to \mathfrak{C}^{\nu+1}$. Then, a hierarchical relational complex Ψ is a special kind of relational persistence complex. An example, in which generalisation increases some Betti numbers is given in Figure 6, where on the left b_0 increases and on the right b_1 is raised by one.*

6 Conclusion

We lay the foundations for relational modelling of topological multi-representation or generalisation by introducing the hierarchical notions of relational chain complexes and their modifications through Euler-Poincaré operators. The strictly topological

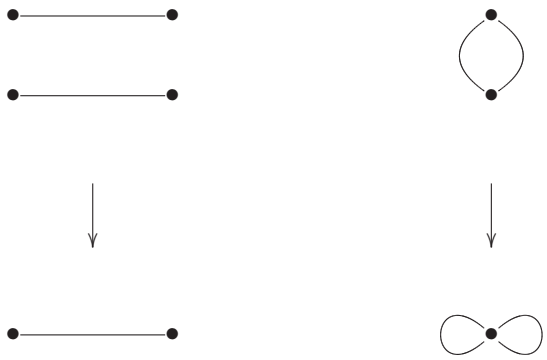

Figure 6: Increasing Betti numbers in generalisation maps. Left: b_0, right: b_1.

point of view allows for relational models as more or less straightforward implementations of their topological kins. The main advantage over the algebraic topological datatype is that the partial matrices allow to query the effective boundary or generalisation of each single cell, loop, shell or other chain, unblurred from the 'sea of zeros' in the sparse total matrices of algebraic topology. The consistency conditions coming from the chain maps yield integrability criteria for the computation of the boundary operators at the different levels from the level of highest detail. This hierarchical point of view opens the door for ultrametric methods in spatial information science.

Acknowledgments

This work is funded by the Deutsche Forschungsgemeinschaft under the project numbers BR 2128/12-1 and BR 3531/3-1. It takes inspiration from the presentation of [9] at the ECCV 2010 conference in Heraklion, Greece. Norbert Paul is thanked for valuable and fruitful discussions.

References

[1] BRADLEY, P. E. From image processing to topological modelling with p-adic numbers. *p-Adic Numbers, Ultrametric Analysis and Applications.* (In press).

[2] BRADLEY, P. E., AND PAUL, N. Comparing G-maps with other topological data structures. Preprint.

[3] BRADLEY, P. E., AND PAUL, N. Using the relational model to capture topological information of spaces. *The Computer Journal 53* (2010), 69–89. DOI: 10.1093/comjnl/bxn054.

[4] BREUNIG, M. *On the way to Component-based 3D/4D geoinformation systems*. LNES 94. Springer, Heidelberg, 2001. DOI: 10.1007/3-540-44950-7.

[5] EGENHOFER, M. H., AND FRANSOZA, R. Point-set topological spatial relations. *Int. J. Geographical Information Science* (1995), 161–174. DOI:10.1080/02693799108927841.

[6] EULER, L. Solutio problematis ad geometriam situs pertinentis. *Comment. acad. scient. Petrop. 8* (1741), 128–140.

[7] LIENHARDT, P. *N*-Dimensional generalized combinatorial maps and cellular quasi-manifolds. *Int. J. Comp. Geom. App. 4* (1994), 275–324. DOI:10.1142/S0218195994000173.

[8] MÄNTYLÄ, M. *An introduction to solid modeling*. Computer Science Pr., 1988.

[9] MUNOZ, D., BAGNELL, J. A., AND HEBERT, M. Stacked hierarchical labeling. In *Proc. ECCV 2010. Part IV* (2010), LNCS 6316, pp. 57–70. DOI: 10.1007/978-3-642-15567-3_5.

[10] ZOMORODIAN, A., AND CARLSSON, G. Computing persistent homology. *Discrete Comput Geom 33* (2005), 249–274. DOI: 10.1007/s00454-004-1146-y.