

Multiresolution Representation of Shapes Based on Cell Complexes

Leila De Floriani, Paola Magillo and Enrico Puppo

Dipartimento di Informatica e Scienze dell'Informazione

Università di Genova

Via Dodecaneso, 35 – 16146 Genova, ITALY

Email: {deflo,magillo,puppo}@disi.unige.it

Abstract. This paper introduces a dimension-independent multiresolution model of a shape, called the Multi-Complex (MC), which is based on decomposition into cells. An MC describes a shape as an initial cell complex approximating it, plus a collection of generic modification patterns to such complex arranged according to a partial order. The partial order is essential to extract variable-resolution shape descriptions in real time. We show how existing multiresolution models reduce to special cases of MCs characterized by specific modification patterns. The MC acts as a unifying framework that is also useful for comparing and evaluating the expressive power of different approaches.

1 Introduction

Multiresolution geometric models support representation and processing of spatial entities at different levels of detail. Such representations have gained recently much of attention in the literature because of their potential impact on applications, such as terrain modeling in geographic information systems, scientific data visualization, virtual reality, etc. The basis for a multiresolution geometric model is the decomposition of the shape it describes into simple elements, called cells.

Cell complexes are used as *discrete models* of a variety of shapes in two, three or higher dimensions. For example, two-dimensional complexes made of polygonal cells are used for describing the boundary of solid objects. In particular, two-dimensional simplicial complexes (triangle meshes) are used for representing surfaces in computer graphics, and for representing terrains in geographic applications. In solid modeling, three-dimensional complexes are used for describing the interior of an object as well as its boundary. d -Dimensional simplicial complexes are used as approximate representations of scalar or vector fields.

The *accuracy* of a cell complex in representing a shape depends on the size, number, and density of its cells: a parameter that we call the *resolution* of the complex. A high resolution, and thus a high number of cells, is needed to produce accurate descriptions. On the other hand, *maximum* accuracy is not always required in *each part* of a shape, but a sufficiently high accuracy for the specific application task can be achieved by *locally adapting* the resolution of a complex in different parts of the shape, thus reducing processing costs and memory space.

The process of building an adaptive cell complex from scratch is usually time-consuming. Sophisticated techniques are required, for instance, to decide where to place the vertices, how to optimize the shape of the cells, how to estimate and bound the approximation error. With a multiresolution approach, most of the time-consuming operations are performed off-line, in order to build a structure which can be queried efficiently on-line, according to the parameters specified by the application task.

Several multiresolution models for cell complexes have been proposed in the literature, but the field still lacks a unifying framework since:

- Most existing models are designed for specific applications, or classes of applications (i.e., virtual reality, rendering, terrain modeling) and designed for certain operations.
- Existing models usually rely on a specific construction techniques: such models can only be obtained from an initial cell complex by applying a *specific type* of transformation operator.
- Most proposed models are a direct abstraction of the data structure used to implement them.

Here, we formally define a general multiresolution model of a shape based on cell complexes, called the *Multi-Complex* (MC), which extends the Multi-Triangulation, proposed in [Pup96,DFPM97]. The MC describes a partially ordered set of modification patterns specifying how an initial coarse cell complex can be iteratively refined in order to obtain more and more accurate representations. The partial order captures dependency relations between modifications that affect the same portion of the complex. This framework permits to select subsets of modification patterns which are consistent with the partial order: when such modification patterns are applied to the initial cell complex, a variety of complexes at different, variable resolutions can be obtained.

The Multi-Complex is independent of the specific type of modification pattern. Existing multiresolution models reduce to special cases of MCs where modification patterns have certain characteristics. In the MC framework, we deal with generic k -dimensional cell complexes in \mathbb{E}^d which decompose a regular shape, and we allow more general modifications, including refining boundaries, changing the topological type, merging and splitting connected components.

The remainder of the paper is organized as follows: In Section 2, we introduce the concept of *modification pattern*, an entity which describes a modification of a cell complex, and plays a fundamental role in the definition of the Multi-Complex. In Section 3, we define the *Multi-Complex* (MC) as a structure which organizes a partially ordered set of modification patterns in the form of a directed acyclic graph and we discuss its more relevant properties. In Section 4, we show how existing multiresolution models can be seen as special instances of a Multi-Complex and we discuss their properties in such a framework. Finally, Section 5 presents some concluding remarks.

2 Preliminaries

2.1 Background Notions

A k -dimensional cell in \mathbb{E}^d is a subset of \mathbb{E}^d homeomorphic to a closed k -dimensional ball. The union as pointset of a set of cells Γ in \mathbb{E}^d is called the *carrier* of Γ . A *cell complex* is a finite set of cells Γ in \mathbb{E}^d of heterogeneous dimension such that their interiors are disjoint, their union cover the carrier, and the boundary of each cell is made of cells of lower dimensions belonging to the complex. A cell γ in a complex Γ is a *top cell* if there is no other cell in Γ bounded by γ . A k -dimensional cell complex is called *regular* if all its top cells are k -dimensional cells.

A special and interesting case is that of simplicial complexes (e.g., triangle and tetrahedral meshes): a k -dimensional *simplex* in \mathbb{E}^d is the locus of the points in \mathbb{E}^d that can be expressed as the convex combination of $k + 1$ affinely independent points; a k -dimensional *simplicial complex* is a cell complex where all top cells are simplices, and every other cell is generated by an appropriate subset of vertices of some top cell.

Thus, the basic elements of an approach to modeling a geometric shape through cell complexes are:

- a *regular cell complex* Γ , where each top cell has an associated *approximation error*, and where, in general, each cell may have application-dependent *attributes* associated with it;
- possibly, a *function* f , piecewise defined on the top cells of Γ .

For instance, $(k + 1)$ -dimensional scalar fields are described by subdividing the k -dimensional domain of the field into a regular k -dimensional simplicial complex, and by associating an approximation of the field to each k -dimensional simplex.

We are interested in modeling at multiple resolutions the cell complex Γ underlying the model of a spatial entity, independently of the function f defined on its top cells and on the attributes of cells.

2.2 Modifications of a Cell Complex

A *modification* is an operation that replaces a cell complex Γ_{old} (possibly contained in a bigger complex Γ) with another cell complex Γ_{new} . The interesting case is when Γ_{old} and Γ_{new} are approximations at two different resolutions of the same (portion of a) shape. A *modification pattern* specifies a modification as the pair of cell complexes $C = (\Gamma_{old}, \Gamma_{new})$. In what follows, we use the symbols $\Pi_1(C)$ and $\Pi_2(C)$ to denote the first and second element, respectively, of the pair of cell complexes in a modification pattern C .

Applying a modification pattern $C = (\Gamma_{old}, \Gamma_{new})$ to a cell complex Γ consists of replacing Γ_{old} with Γ_{new} in Γ . The result of applying C to Γ , denoted with $\Gamma[C]$, is the set of cells $(\Gamma \setminus \Gamma_{old}) \cup \Gamma_{new}$. Note that set $\Gamma[C]$ may not always be

a cell complex. We are interested in the case where $\Gamma[C]$ is a cell complex, and C does not open or close holes in the boundary of Γ . This corresponds to the notion of *validity* of a modification pattern for a cell complex:

Definition 1. Let Γ be a regular k -dimensional cell complex in \mathbb{E}^d and $C = (\Gamma_{old}, \Gamma_{new})$ be a k -dimensional modification pattern in \mathbb{E}^d . C is a valid modification pattern for Γ if and only if

1. $\Gamma_{old} \subseteq \Gamma$ (as sets of cells);
2. the cell complex obtained by deleting Γ_{old} from Γ does not intersect Γ_{new} outside the common cells of the boundaries of Γ_{old} and Γ_{new} ;
3. the boundary of Γ_{new} must match with the boundary of the “hole” created in Γ when removing Γ_{old} .




Fig. 1. Six modification patterns on a two-dimensional cell complex: C_1 and C_2 are valid; C_3 is not valid since it violates condition 2; C_4 is not valid since it violates condition 3.

Examples of valid and invalid modification patterns are shown in Figure 1. It has been proven that the result $\Gamma[C]$ of applying a valid k -dimensional modification pattern C to a regular k -cell complex Γ is either empty, or is a regular k -dimensional cell complex [Mag98].

For the purpose of defining a multiresolution model, we are interested in k -dimensional *refinement patterns*, i.e., modification patterns where Γ_{new} contains more k -cells than Γ_{old} , and in *minimal* modification patterns, i.e., modification patterns that cannot be split into two or more valid modifications to be performed in a sequence. Modification patterns that induce changes in the topological type of the shape (e.g., increasing or decreasing its genus, merging or splitting connected components) will be allowed.

2.3 Sequences of Modification Patterns

A sequence of modification patterns is *valid* for a cell complex Γ_0 if and only if every modification pattern C_i in the sequence is valid for the complex obtained

as the result of successively applying to Γ_0 all the patterns preceding C_i in the sequence.

We say that a modification pattern C' *directly depends on* another modification pattern C (and, thus, C *directly blocks* C') if and only if the effect of applying C' is removing some of the cells introduced by applying C . In a set of modification patterns, two modification patterns depend on each other if they are in the transitive closure of the relation of direct dependency; in a similar way we can define when they block each other.

In a generic valid sequence a cell γ may be created and removed several times and this creates cycles in the relation of dependency. Thus, we define a modification pattern as *non-redundant* with respect to a set of modification patterns if it does not recreate k -cells eliminated by some other modification pattern of the set. It is easy to verify that, in non-redundant and valid sequences, each cell either is in the initial complex, or there is exactly one modification pattern specifying its “creation”; moreover, each cell either is in the final complex, or there is exactly one modification pattern specifying its “deletion”.

Even if two modification patterns are independent, they cannot necessarily be applied to the same cell complex without interfering. We say that C_i and C_j are not *in conflict* when all the pairs of complexes $\Pi_1(C_i)$ and $\Pi_1(C_j)$, $\Pi_1(C_i)$ and $\Pi_2(C_j)$, $\Pi_2(C_i)$ and $\Pi_1(C_j)$, and $\Pi_2(C_i)$ and $\Pi_2(C_j)$ intersect at most in a subset of their boundary cells, which are preserved in the modifications represented by C_i and C_j . This means that any of such pairs share the same subset of boundary cells, and the union of all their cells forms a complex. A set of modification patterns is thus *conflict-free* if and only if there are no conflicts between pairs modification patterns that are independent.

3 The Multi-Complex and its Properties

In a Multi-Complex, a *partially ordered set of modification patterns* of a cell complex is encoded in the form of a *directed acyclic graph* (DAG), where nodes represent modification patterns, and arcs represent relations of direct dependency between them. An MC is defined by abstracting over the relation of direct dependency from a *valid, non-redundant, and conflict-free* sequence of *minimal refinement* patterns. The non-redundancy of the sequence ensures that the direct dependency relation is a partial order. Since all nodes describe refinement patterns, the resolution monotonically increases while traversing the DAG according to the direction of its arcs. The validity and absence of conflicts ensure that, for every subset of modification patterns “closed” with respect to the dependency relation, the result of applying its elements, sorted in any total order consistent with the partial one, is a cell complex.

Definition 2. Let Γ_0 be a k -dimensional cell complex, C_1, C_2, \dots, C_n be a sequence of k -dimensional modification patterns such that:

- for every $1 \leq i \leq n$, C_i is a refinement pattern, it is minimal, and $\Pi_1(C_i) \neq \emptyset$;
- C_1, C_2, \dots, C_n is a valid sequence for Γ_0 ;

- sequence $C_0, C_1, C_2, \dots, C_n$ is non-redundant, where $C_0 = (\emptyset, \Gamma_0)$;
- set $\{C_1, \dots, C_n\}$ is conflict-free.

Then, a Multi-Complex (MC) is a directed graph \mathcal{M} where:

- \mathcal{M} has n nodes, which are in one-to-one correspondence with the modification patterns of set $\{C_0, \dots, C_n\}$;
- there is an arc (C_i, C_j) if and only if C_j directly depends on C_i .

A simplicial Multi-Complex is a special case of MC where each modification pattern consists of a pair of simplicial complexes. The MC of Figure 2(a) is a simplicial MC.

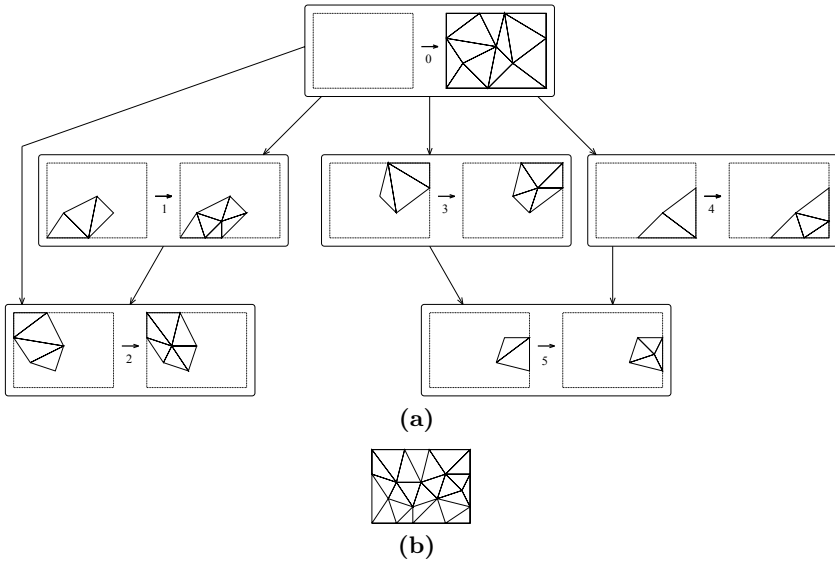


Fig. 2. (a) A Multi-Complex: each node is a modification pattern C_i , where $\Pi_1(C_i)$, and $\Pi_2(C_i)$ are the complex on the left side and on the right side, respectively; (b) its front cell complex.

Given a Multi-Complex \mathcal{M} , and a consistent sequence of its nodes (including the root C_0) C_0, C_1, \dots, C_n , the *front cell complex* of \mathcal{M} is

$$\mathcal{F}(\mathcal{M}) \equiv \Pi_2(C_0)[C_1] \cdots [C_n].$$

The front cell complex is a regular cell complex, and it is uniquely defined because it is independent on the specific sequence considered [Mag98]. Figure 2(b) shows the front complex associated with the MC of Figure 2(a).

A sub-MC \mathcal{M}' of a Multi-Complex \mathcal{M} identifies a subset of the modification patterns of \mathcal{M} which contains the root C_0 and is closed with respect to the dependency relation (i.e., it contains the parents of any of its nodes). Since the

MC-nodes represent refinement patterns, $\mathcal{F}(\mathcal{M}')$ is a cell complex at a lower resolution than the front cell complex of \mathcal{M} (see Figure 3). Different sub-MCs of \mathcal{M} provide cell complexes at different resolutions. Cell complexes at *intermediate resolutions* are front complexes of sub-MCs.




Fig. 3. A sub-MC containing nodes $\{0, 3, 4, 5\}$ (a), and its front cell complex (b).

In what follows, we will discuss some properties of an MC which are relevant to the efficiency of its encoding data structures, and of the query algorithms operating on it, as discussed in [DFMP98, Mag98]. Such properties are determined by the construction algorithm used to produce the initial cell complex and the sequence of modification patterns which form the basis of a Multi-Complex.

We define the *height* of a Multi-Complex as the maximum number of arcs in a path from the root to a leaf; the *width* of a node C as the number of top cells in $\Pi_2(C)$, and the width of a Multi-Complex as the maximum width of its nodes; the *size* of a Multi-Complex as the total number of top cells in it, and the *size* of the front cell complex of an MC as the number of its top cells.

We are interested in Multi-Complexes which have a height logarithmic in the number of its top cells, and a width bounded from above by a small constant b . In [DFMP97, Mag98], we have shown that bounded width and logarithmic height are important for the efficiency of traversal operations on an MC required in query processing, such as a point location or a range query.

We say that an MC \mathcal{M} has a *linear growth* if and only if there exists a positive real number b such that for every sub-MC \mathcal{M}' of \mathcal{M} , the ratio between the size of \mathcal{M}' and the size of the front cell complex of \mathcal{M}' is less or equal to b . It can be easily seen that if an MC has bounded width, then it has a linear growth as well.

If an MC has a linear growth, then the size of any sub-MC is linear in the size of its front complex. In particular, the size of the MC itself is linear in the size of its front cell complex. This means that the multiresolution structure introduces only a linear storage overhead with respect to a simple cell complex at the maximum resolution.

In [Pup96,Mag98], we have shown that a linear growth is fundamental to achieve optimal time complexity (i.e., linear in the output size) for algorithms which extract a cell complex at a given resolution from an MC: since these algorithms perform a DAG traversal to find the appropriate sub-MC, linear growth guarantees that the size of the traversed sub-MC is linear in the output size.

4 Existing Multiresolution Models as Multi-Complexes

In this section, we show how multiresolution models presented in the literature can be seen as special cases of Multi-Complexes. Each of these models gives a class of MCs characterized by certain properties of the modification patterns, and, consequently, by certain properties in terms of width, height and growth. Moreover, each existing model provides an implementation of an MC with a specific data structure.

Often some information present in the MC is either lost or incorrectly represented in some data structures. Examples of lost information are found in models which encode a linear sequence [Hop96,KS97] instead of the DAG; we will see examples of incorrectly represented information, for instance, in nested models. Such an aspect will be examined for existing models in Sections 4.1, and 4.2. For the sake of brevity, we explain the interpretation of just two-dimensional multiresolution models as two-dimensional MCs. Higher-dimensional models can be interpreted in a similar way.

4.1 Nested Models

Nested models are based on a nested subdivision of a square or triangular domain into scaled copies of it and are usually described by a tree, where each node corresponds to a complex which covers one of the cells of its parent node. Examples of nested models are the quadtree, the quaternary triangulation, the restricted quadtree, the hierarchies of right triangles and the hierarchical irregular triangulations. Figure 4 (a) shows an example of a hierarchical irregular triangulation.

In a k -dimensional nested model, a node represents the refinement of a top-cell γ of its parent node into a complex I' . The corresponding modification

pattern is $C = (\{\gamma\}, \Gamma')$, where Γ is the complex contained in the parent node. In general (and in most existing nested models), C is not a valid modification pattern for Γ , since Γ' may split some of the $(k-1)$ -facets of γ : the intersections of the cells of $\Gamma \setminus \{\gamma\}$ with those of Γ' do not coincide with the common cells of their boundary complexes, thus violating the second condition in the definition of Multi-Complex. Figures 4 (a) and (b) illustrate this problem. A consistent cell complex is especially important when interpolating functions are defined over the cells, as in the case of scalar and vector fields.




Fig. 4. (a) The tree representing a nested model. (b) The complex Γ' refining a top-cell γ does not define a valid modification pattern for the complex Γ containing γ ; in fact the boundary of Γ' does not match with that of the adjacent cell to γ in Γ . (c) Clusters of tree nodes that become MC-nodes when interpreting the nested model as an MC.

If Γ' splits a $(k-1)$ -facet γ' of γ , this means that γ can be replaced consistently with Γ' only if also the cell γ_1 , adjacent to γ along γ' in Γ , is replaced by a complex Γ'_1 splitting facet γ' in the same way. Thus, nodes in the tree representing the refinement of γ and γ' must be clustered to form one node of the Multi-Complex. The result of such node clustering is shown in Figure 4 (c).

A Multi-Complex representing a k -dimensional nested model is obtained as follows:

- the MC-nodes are obtained from the nodes of the tree by iteratively joining pairs of nodes corresponding to complexes in which two adjacent k -cells have been refined by decomposing their common $(k-1)$ -facet; this defines a mapping of nodes of the tree to MC-nodes.
- the MC arcs are obtained from the tree arcs according to the following rule: for every arc (N_1, N_2) of the tree, if N_1 maps to an MC-node C_1 and N_2 maps to an MC-node C_2 , then arc (C_1, C_2) represents (N_1, N_2) in the MC; this defines a mapping of the arcs of the tree to arcs of the MC.

The Multi-Complex derived from a given nested model shows what consistent cell complexes can be obtained from it (i.e., cell complexes that avoid cracks in the graph of a scalar or vector field represented by the nested model). Such

consistent cell complexes are in one-to-one correspondence with the sub-MCs of the corresponding Multi-Complex. The range of (possibly inconsistent) sets of cells obtained from a nested model is in one-to-one correspondence with the analogous sub-structures of the tree. Since several nodes of the tree are merged into one MC-node, it follows that some sub-structures of the tree are not captured by the MC: these are exactly those cuts leading to inconsistent representations.

Since no knowledge about different tree nodes that must be clustered to form an MC-node is explicitly encoded in the tree, algorithms for extracting generic (i.e., variable resolution) complexes from nested models are more involved.

Nested models based on the recursive subdivision of a domain according to a fixed pattern that splits *every* $(k - 1)$ -facet of the current k -cell (e.g., quadrees and their extensions [Sam90], quaternary triangulations [GG79]), correspond to Multi-Complexes where the DAG reduces to a list: all the nodes in a level of the tree must be merged together into one MC-node. Thus, the range of consistent complexes that can be obtained from such models reduces to one for every tree level.

Nested models that leave some $(k - 1)$ -facet of a k -cell γ unrefined when expanding γ into a complex (e.g., hierarchical triangulations [SP92,DFP95], hierarchies of right triangles [DWS⁺97,LKR⁺96]) may have a better expressive power. In general, if $(k - 1)$ -cells are split too often, then the nodes of the tree tend to be clustered into large MC-nodes, hence exhibiting a low expressive power. On the other hand, the persistence of $(k - 1)$ -cells across the levels of the tree tends to produce slivery cells, which are not desirable. In two dimensions, the best compromise seems to have been achieved with hierarchies of right triangles, where, at each level, one edge of each triangle is split, and two edges are retained to be split at the next level (see Figure 5).

4.2 Evolutionary Models

Evolutionary models track the evolution of an initial cell complex Γ_0 through a sequence of local modifications, which is assumed to be a valid sequence for Γ_0 .

A Multi-Complex for such models represents the local modifications of the sequence in its nodes, and encodes their relations of direct dependency in its arcs: there is an arc from a node $C_i = (\Gamma_i, \Gamma'_i)$ to a node $C_j = (\Gamma_j, \Gamma'_j)$ if and only if the intersection of Γ'_i and Γ_j (as sets of cells) contains at least one top-cell. The validity and non-redundancy of the sequence ensures that the resulting graph is a DAG.

In most existing models, the local modifications are performed by applying a specific operator (see Figure 6):

- In [Hop96,XESV97,Hop97,MMS97,GTLH98,KCVS98] the refinement pattern is a vertex split, which expands a vertex p into an edge, two edges incident at p into two triangles, and warps the triangles surrounding p accordingly.
- In [DF89,CDFM⁺94,dBD95,BFM95,KS97,CPS97,DFMP98] the refinement pattern consists of inserting a new vertex p , and replacing a subset of triangles in the neighborhood of p with a set of triangles incident at p . Note that




Fig. 5. A hierarchy of right triangles (a), with meshes shown level by level, and the clusters of tree nodes corresponding the nodes of the Multi-Complex (b).




Fig. 6. Vertex insertion and vertex split. The shaded triangles are those involved in the modification pattern.

a vertex split where the edge created by splitting p has p itself as one of its endpoints can be regarded as a special case of vertex insertion.

In all models, the local modifications guarantee the non-redundancy and the minimality of the modifications.

Evolutionary models proposed in the literature represent the local modifications and/or the cells involved within them by using different structures, from sequential lists to DAGs. In the following we illustrate the relationship between those structures and the corresponding Multi-Complex.

Sequential Models Sequential models [KS97,Hop96] keep the local modifications in a totally ordered list, which is consistent with the partial order represented in the Multi-Complex. Usually, such total order is either the same, or the reverse of the order in which the local modifications have been performed during the construction of the model. Dependencies between local modifications are not




Fig. 7. (a) A sequence of edge collapses, to be read reversed as a sequence of vertex splits. Shaded triangles correspond to the two meshes forming a modification pattern in the MC. Triangles marked with a circle are those considered in [Hop97] for determining the vertex dependencies. (b) The corresponding binary forest of vertices.

explicitly stored, and thus they must be reconstructed on-line when extracting a complex at variable resolution by performing considerable backtricking.

In terms of the corresponding MC, the models in [KS97] and [Hop96] do not guarantee bounded width, logarithmic height or linear growth.

Hierarchies of Vertices Hierarchies of vertices identify a local modification (namely, a vertex split) with the vertex v being split. Thus, dependencies between modification patterns become dependencies between vertices. Most models (e.g., [Hop97,XESV97]) distinguish between two cases of dependency:

1. if a vertex v is an endpoint of the edge created by splitting w , then v depends on w ;
2. v also depends on (some of) the vertices that are incident at v when v is split.

Dependencies of type 1 are stored as a binary tree of vertices (every vertex split creates two new vertices, which become its children); dependencies of type 2 are

Multi-Complex

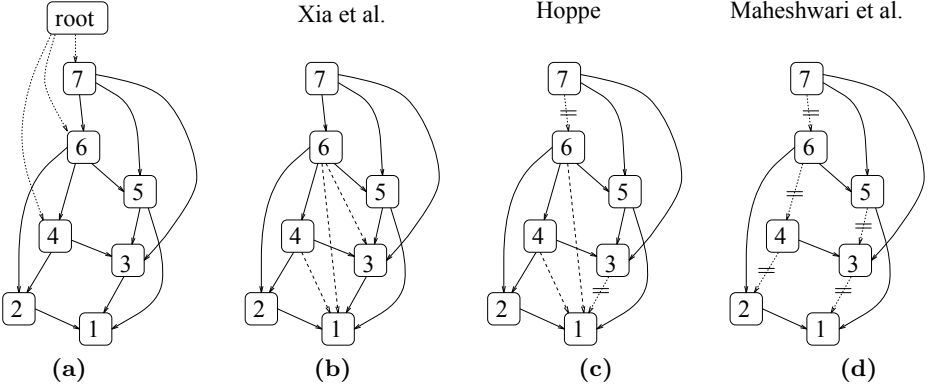


Fig. 8. (a) The Multi-Complex corresponding to the sequence of edge splits in Figure 7; the arcs emanating from the root are shown dashed since the root does not represent a vertex splits. (b) The DAG representing vertex dependencies according to [XESV97]; (c) The DAG representing vertex dependencies according to [Hop97]; (d) The DAG representing vertex dependencies according to [MMS97]; redundant arcs are dashed, missing arcs are dotted and marked with ‘=’.




Fig. 9. Edge collapse and its reverse vertex split. According to [XESV97], vertices v_1 and v_2 depend on the vertices of all triangles incident at v (i.e., $w_1, w_2, p_1, p_2, p_3, p_4, p_5$); according to [Hop97], v_1 and v_2 depend on the vertices of the shaded triangles (i.e., $w_1, w_2, p_1, p_3, p_4, p_5$); according to [MMS97], v_1 and v_2 depend just on w_1 and w_2 .

stored separately. The various authors define dependencies of type 2 in different ways.

Xia, El-Sana, and Varshney [XESV97] consider a vertex v as depending on all the vertices that are adjacent to v when v is split. With respect to the Multi-Complex, the resulting DAG contains redundant arcs, which correspond to paths in the MC (see Figure 8 (b)).

Hoppe [Hop97] considers a vertex v as dependent only of the vertices of the four facets incident at the two edges that are expanded into triangles when v is split (see Figure 9). However, this does not capture all the existing dependencies, and there are still redundant arcs (see Figure 8 (c)).

Progressive TINs [MMS97] consider a vertex p as dependent of its parent in the binary tree, and of the other endpoints of the two edges that are expanded

into triangles when splitting p . Again, some dependencies are missing (see Figure 8 (d)).

The model by Luebke and Erikson [LE97] is based on modifications that are “generalized” vertex splits: a vertex may be split into two vertices not necessarily connected by a segment (such an operation is the inverse of not only edge collapse but also vertex-pair collapse, when the two collapsed vertices are not adjacent in the mesh). This model keeps just the binary tree of vertices, and ignores all other existing dependencies.

The problem of redundant arcs is intrinsic to the binary tree of vertices: a vertex split may depend only indirectly on its parent, since other modifications may have altered the configuration of triangles between the parent and the child. For instance, the binary tree shown in Figure 7 contains arc (6, 1), which is not in the corresponding MC shown in Figure 8.

Extra arcs increase the space requirements of data structures but do not compromise the integrity of the model. On the contrary, missing arcs are a problem because DAG traversal algorithms working on such models may extract inconsistent cell complexes.

A cell complex extracted from a model affected by missing dependency links [Hop97,LE97,MMS97] may contain intersecting triangles, and triangles which were do not belong to any modification pattern in the model. The problem of “extra” triangles is relevant in applications where it is necessary to know the approximation error of any triangle with respect to the original shape.

Guéziec et al. [GTLH98] store vertex dependencies explicitly by means of a DAG, but they consider dependencies as in [XESV97], thus their DAG contains some redundant arcs.

The models of [Hop97,GTLH98,MMS97] do not guarantee bounded width, logarithmic height or linear growth. The Multi-Complex corresponding to the model of [XESV97] has linear growth, and logarithmic height, but it does not necessarily have bounded width.

5 Concluding Remarks

In this paper, we have defined the Multi-Complex (MC), a general model for multiresolution cell complexes. An MC consists of a DAG that represents a partially ordered set of modifications refining an initial cell complex; the partial order is induced by the relation of direct dependency, where a modification that deletes some top-cells from the current complex directly depends on those modifications that have previously introduced such top-cells.

Cell complexes at variable resolution can be obtained from an MC by selecting a subset of modifications consistent with the partial order.

We have shown that the Multi-Complex encompasses all existing multiresolution models and helps understanding the range of cell complexes at variable resolution that a model can provide. Thus, the MC acts as a unifying framework that is useful for comparing the expressive power of different approaches.

The definition of the Multi-Complex, and of most other multiresolution models developed in the literature is limited to regular cell complexes. However, in some applications it is important to deal with non-regular shapes, and modification patterns that change the dimensionality of (portions of) a shape. Extensions of the MC framework in such a perspective will be the subject of future research.

Acknowledgments

Part of the work described in this paper has been developed while the first author was on leave from the University of Genova at the University of Maryland Institute for Applied Computer Studies (UMIACS). The support of National Science Foundation (NSF) Grant “The Grand Challenge” under contract BIR9318183 is gratefully acknowledged.

References

- BFM95. M. Bertolotto, L. De Floriani, and P. Marzano. Pyramidal simplicial complexes. In *Proceedings 4th International Symposium on Solid Modeling*, pages 153–162, Salt Lake City, Utah, U.S.A., May 17-19 1995. ACM Press.
- CDFM⁺94. P. Cignoni, L. De Floriani, C. Montani, E. Puppo, and R. Scopigno. Multiresolution modeling and rendering of volume data based on simplicial complexes. In *Proceedings of 1994 Symposium on Volume Visualization*, pages 19–26. ACM Press, October 17-18 1994.
- CPS97. P. Cignoni, E. Puppo, and R. Scopigno. Representation and visualization of terrain surfaces at variable resolution. *The Visual Computer*, 13:199–217, 1997. (A preliminary version appeared on “Scientific Visualization ’95”, Proc. Inter. Symposium, World Scientific, pp.50-68).
- dBD95. M. de Berg and K. Dobrindt. On levels of detail in terrains. In *Proceedings 11th ACM Symposium on Computational Geometry*, pages C26–C27, Vancouver (Canada), 1995. ACM Press.
- DF89. L. De Floriani. A pyramidal data structure for triangle-based surface description. *IEEE Computer Graphics and Applications*, 8(2):67–78, 1989.
- DFMP97. L. De Floriani, P. Magillo, and E. Puppo. Building and traversing a surface at variable resolution. In *Proceedings IEEE Visualization 97*, pages 103–110, Phoenix, AZ (USA), October 1997.
- DFMP98. L. De Floriani, P. Magillo, and E. Puppo. Efficient implementation of multi-triangulations. In *Proceedings IEEE Visualization 98*, Research Triangle Park, NC (USA), October 1998.
- DFP95. L. De Floriani and E. Puppo. Hierarchical triangulation for multiresolution surface description. *ACM Transactions on Graphics*, 14(4):363–411, October 1995.
- DFPM97. L. De Floriani, E. Puppo, and P. Magillo. A formal approach to multiresolution modeling. In R. Klein, W. Straßer, and R. Rau, editors, *Geometric Modeling: Theory and Practice*. Springer-Verlag, 1997.
- DWS⁺97. M. Duchaineau, M. Wolinsky, D.E. Siget, M.C. Miller, C. Aldrich, and M.B. Mineed-Weinstein. Roaming terrain: Real-time optimally adapting meshes. In *Proceedings IEEE Visualization’97*, pages 81–88, 1997.

- GG79. D. Gomez and A. Guzman. Digital model for three-dimensional surface representation. *Geo-Processing*, 1:53–70, 1979.
- GTLH98. A. Guéziec, G. Taubin, F. Lazarus, and W. Horn. Simplicial maps for progressive transmission of polygonal surfaces. In *Proceeding ACM VRML98*, pages 25–31, 1998.
- Hop96. H. Hoppe. Progressive meshes. In *ACM Computer Graphics Proc., Annual Conference Series, (SIGGRAPH '96)*, pages 99–108, 1996.
- Hop97. Hugues Hoppe. View-dependent refinement of progressive meshes. In *ACM Computer Graphics Proc., Annual Conference Series, (SIGGRAPH '97)*, pages 189–198, 1997.
- KCVS98. L. Kobbelt, S. Campagna, J. Vorsatz, and H.P. Seidel. Interactive multi-resolution modeling of arbitrary meshes. In *Comp. Graph. Proc., Annual Conf. Series (SIGGRAPH '98)*, ACM Press, 1998. to appear.
- KS97. R. Klein and W. Straßer. Generation of multiresolution models from CAD data for real time rendering. In R. Klein, W. Straßer, and R. Rau, editors, *Theory and Practice of Geometric Modeling (Blaubeuren II)*. Springer-Verlag, 1997.
- LE97. D. Luebke and C. Erikson. View-dependent simplification of arbitrary polygonal environments. In *ACM Computer Graphics Proc., Annual Conference Series, (SIGGRAPH '97)*, 1997. (to appear).
- LKR⁺96. P. Lindstrom, D. Koller, W. Ribarsky, L.F. Hodges, N. Faust, and G.A. Turner. Real-time, continuous level of detail rendering of height fields. In *Comp. Graph. Proc., Annual Conf. Series (SIGGRAPH '96)*, ACM Press, pages 109–118, New Orleans, LA, USA, Aug. 6-8 1996.
- Mag98. Paola Magillo. *Spatial Operations on Multiresolution Cell Complexes*. PhD thesis, Dept. of Computer and Information Sciences, U. of Genova, 1998.
- MMS97. A. Maheshwari, P. Morin, and J.-R. Sack. Progressive TINs: Algorithms and applications. In *Proceedings 5th ACM Workshop on Advances in Geographic Information Systems, Las Vegas*, 1997.
- Pup96. E. Puppo. Variable resolution terrain surfaces. In *Proceedings Eight Canadian Conference on Computational Geometry*, pages 202–210, Ottawa, Canada, August 12-15 1996. Extended version to appear under title “Variable resolution triangulations” in *Computational Geometry Theory and Applications*.
- Sam90. H. Samet. *Applications of Spatial Data Structures*. Addison Wesley, Reading, MA, 1990.
- SP92. Lori Scarlatos and Theo Pavlidis. Hierarchical triangulation using cartographics coherence. *CVGIP: Graphical Models and Image Processing*, 54(2):147–161, March 1992.
- XESV97. J.C. Xia, J. El-Sana, and A. Varshney. Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):171–183, 1997.