

A Generic Programming Approach to Multiresolution Spatial Decompositions

Vinícius Mello¹, Luiz Velho¹, Paulo Roma Cavalcanti² and Cláudio Silva³

¹IMPA – Instituto de Matemática Pura e Aplicada
{vinicius,lvelho}@visgraf.impa.br

²Federal University of Rio de Janeiro - UFRJ
roma@lcg.ufrj.br

³AT&T Labs
csilva@research.att.com

Abstract

We present a generic programming approach to the implementation of multiresolution spatial decompositions. From a set of simple and necessary requirements, we arrive at the Binary Multitriangulation (BMT) concept. We also describe a data structure that models the BMT concept in its full generality. Finally, we discuss applications of the BMT to visualization of volumetric datasets.

1 Introduction

Generic programming was born from the observation that most algorithms rely on a few basic semantic assumptions about the data structures, and not on any particular implementation of these structures. Given a problem, the generic programming basic task is to isolate these essential concepts, framing them in a well defined interface where semantic requirements and computational complexity guarantees are clearly posed. The algorithms that comply with that interface are free from idiosyncrasies of data structures, which can be changed or even replaced by procedural schemes. The great success of C++ Standard Template Library is the main proof on behalf of that methodology [29].

The main contribution of this paper is to show how generic programming techniques can be used to build a computational framework to deal with multiresolution spatial decompositions. We have studied from combinatorial topology classics like [1] to modern works on multiresolution modeling [12] in order to identify the meaningful concepts. As a result, we arrive at a new concept called Binary Multitriangulation (BMT) that is a particular case of the Multiresolution Simplicial Model (MSM) described in [12], but more manageable and closer in spirit to well established procedures of combinatorial topology. The BMT concept can also be regarded as a 3-dimensional extension of variable resolution structures like [31].

Currently, generic programming methodology is being used as design philosophy of many libraries in several areas: computational geometry (CGAL [10]), combinatorics (BGL [26]) and scientific computing (MTL [27]), for instance. Concerning to the specific problem of spatial decompositions, our work resembles the GrAL library [2], although we were unaware of that until just before the publication of this paper. The similarities can be ascribed to the application of the same paradigm (generic programming) to a same problem (spatial decompositions) with the same search for conceptual rigor. More important, nonetheless, are the dissimilarities: while the GrAL library concentrates in fixed resolution cell decompositions, we focus on multiresolution simplicial decompositions.

A common point of all those libraries is the omnipresence of C++ language, the most fitted language to generic programming (specially after the recent ISO/ANSI standardization). This fact justifies the extensive use of C++ code in this paper, but we must emphasize that there is no intrinsic dependence to specific programming languages in the concepts we describe. In a sense, genericity can be achieved in almost any language, with more or less effort. We even can regard the pioneering work of Guibas and Stolfi [14], Laszlo and Dobkin [7] and Mäntylä [18] as first attempts to attain genericity in surface and solid modeling. The C++ language just provides a set of built in facilities to do it.

The paper is organized as follows. In section 2, we set the context where multiresolution spatial decompositions are needed and discuss informally some of their advantages. Section 3 is dedicated to the detailed examination of the concepts we isolate, in increasing order of complexity. In section 4, we describe data structures that models the previously defined concepts. Some applications are presented in section 5. Finally, section 6 contains concluding remarks and indications of future works.

2 Background

If multiresolution methods are significant in the processing of triangle meshes [9, 8], they are indispensable in the case of tetrahedra meshes, or spatial decompositions, as we prefer to call it, since the complexity of the mesh increases with the power of its dimension. Therefore, most applications dealing with three-dimensional data, like scientific visualization, medical imaging, geoprocessing etc., will benefit of techniques that allow the user to extract from the original data an equivalent representation, in a sense will be made precise subsequently, but in a resolution more adequate to the task at hand. Because the user often doesn't know *a priori* the desired resolution, the solution is to store a good set of possible resolutions and to give him the ability of browsing between them. That is the essence of the multiresolution methods.

A cost-benefit analysis of multiresolution methods with respect to the current technology is presented in [4]. The main conclusion of this analysis is that, in the case of direct volume rendering applications (DVR), the graphics constraints are stronger than memory constraints. In fact, one can store a mesh three orders of magnitude larger than that one can visualize with DVR. It is reasonable to assume that similar conclusions can be extended to other applications dealing with volume data. Thus, the extra cost to store a multiresolution mesh structure is compensated by the flexibility to choose the most adequate resolution.

It remains to define more precisely what is a tetrahedral mesh. Although the basic intuitive concept is clear, it is not so clear that some properties of a mesh are independent of the geometry of its constituent tetrahedra or, more specifically, of their spatial embedding. The area of mathematics that studies such properties is called *combinatorial topology*. A *manifold* is a well known mathematical object and it is very useful in combinatorial topology. A combinatorial manifold is characterized by a certain uniformity in the way its parts are glued. Despite of the fact that in some applications it is necessary to consider “non-manifold” structures, the concept of a manifold is sufficient for most applications.

In the next session, we will describe a series of concepts progressively more complex, until we arrive at a definition of a combinatorial manifold. Each concept will be followed by an API which defines the operations required to work with the concept. We will postpone the introduction of geometric concepts as much as possible, in order to clearly isolate the topological properties.

However, geometric concepts, such as volume, area, aspect ratio, etc., are of fundamental importance in applications and are deeply related with the mechanisms that are employed to select a particular mesh resolution from a multiresolution structure. In general, we deal with functions whose domain is the combinatorial manifold, and we would like to have a more refined mesh in regions where these functions exhibit high variations. This property is called *adaptivity*. We intend to discuss techniques for generation and processing of multiresolution adaptive tetrahedral meshes in a future paper.

3 Concepts

In this section, we adopt the following strategy in describing the concepts: we will define the mathematical objects involved, the name and type signature of the requirements, and let the semantics be derived from them and from hints in the text body. Auxiliary tables containing associated types and notations fulfill the description. Concerning to type signatures, each concept has a *trait class* where all type information is encoded. Trait classes are essentially a mechanism to ensure algorithm independence of data structure implementation [30]. In order to clarify the programs presentation to the noninitiated in the C++

description	type
Vertex descriptor	as3c_traits<T>::vertex_descriptor
Edge descriptor	as3c_traits<T>::edge_descriptor
Face descriptor	as3c_traits<T>::face_descriptor
Simplex descriptor	as3c_traits<T>::simplex_descriptor
Complex vertices iterator	as3c_traits<T>::vertex_iterator
Complex edges iterator	as3c_traits<T>::edge_iterator
Complex faces iterator	as3c_traits<T>::face_iterator
Complex simplices iterator	as3c_traits<T>::simplex_iterator

Table 1: Associated types of an AS3C. T is a type that models an AS3C. The `*_descriptor` types are intended to mean “small types”, that is, types which objects can be passed by value without overhead. The `*_iterator` types must be at least forward iterators.

technicalities, we will apply the notation defined in the auxiliary tables to the sample code and we will omit the `template` clause (see program 1 and program 2).

Some definitions bellow differs from the usual (see [6], for instance), but this happens because we have choosen equivalent definitions easily translatable to algorithm requirements.

3.1 Abstract Simplicial 3-Complex

Definition 1. *Given a finite set V , called vertex set, an abstract simplicial complex on V is a set K of subsets of V verifying the following properties:*

1. For each $\nu \in V$, $\{\nu\} \in K$;
2. If $\sigma \in K$ and $\phi \subset \sigma$, then $\phi \in K$. An element σ of K is called *simplex* and the subsets of σ are called *faces*;
3. There is a total ordering on the vertices of each simplex of K such that the ordering on the vertices on any face of a simplex σ is the ordering induced from the ordering on the vertices of σ .

If $n + 1$ is the cardinality of a simplex $\sigma \in K$, we say that σ is a n -simplex and K is an *abstract simplicial n -complex* if the largest simplex of K is a n -simplex. We will use the same name *simplex* to mean the subcomplex of K formed by the faces of a simplex σ , and we define $\partial\sigma = \{\tau \in K : \tau \subsetneq \sigma\}$, that is, the *boundary* of σ .

We are interested here in abstract simplicial 3-complexes, AS3C for short. In this case, we adopt the terminology *vertex*, *edge*, *face* and *simplex* for 0, 1, 2 and 3-simplex, respectively.

Item 3 from definition 1 has a twofold propose: it provides a “canonical form” for each simplex and enables us to define the *face operator* d_i , that assigns for each simplex σ the face of σ obtained by removing the i -th vertex. The face operator satisfies

$$d_i d_j = d_{j-1} d_i, \text{ if } i < j. \tag{1}$$

The existence of operators satisfying (1) is sufficient to recover all relations between the faces of a simplex¹. Figure 1 shows all those relations. We exploit this fact to define a minimum set of requirements on an AS3C (see table 3).

¹Face operators appear in algebraic topology in the definition of simplicial sets [19].

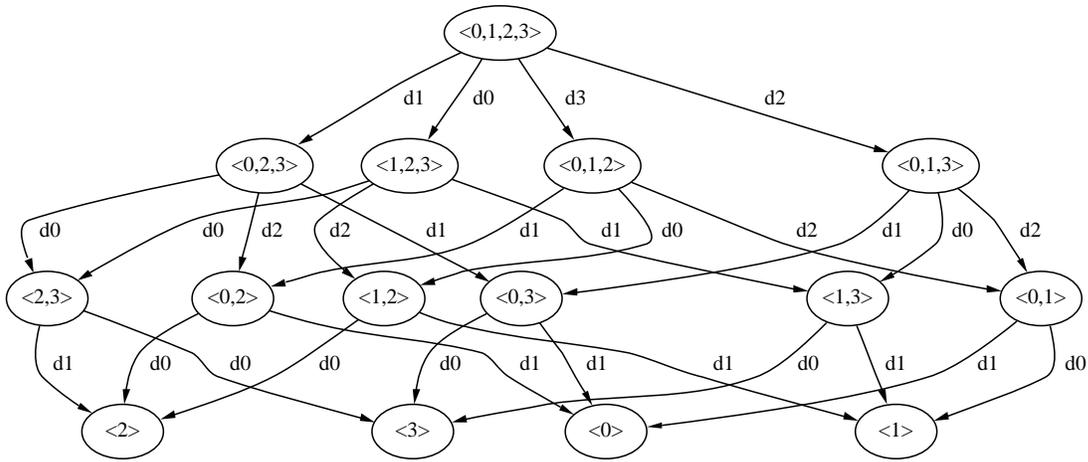


Figure 1: Face operator graph. The nodes are the simplex's faces and the edges are the face operators d_i .

symbol	definition
vertex	<code>typedef as3c_traits<T>::vertex_descriptor vertex;</code>
v	a object of type vertex
edge	<code>typedef as3c_traits<T>::edge_descriptor edge;</code>
e	a object of type edge
face	<code>typedef as3c_traits<T>::face_descriptor face;</code>
f	a object of type face
simplex	<code>typedef as3c_traits<T>::simplex_descriptor simplex;</code>
s	a object of type simplex
vi	<code>typedef as3c_traits<T>::vertex_iterator vi;</code>
ei	<code>typedef as3c_traits<T>::edge_iterator ei;</code>
fi	<code>typedef as3c_traits<T>::face_iterator fi;</code>
si	<code>typedef as3c_traits<T>::simplex_iterator si;</code>

Table 2: AS3C related notation.

expression	return type
<code>empty_vertex(t)</code>	<code>vertex</code>
<code>empty_edge(t)</code>	<code>edge</code>
<code>empty_face(t)</code>	<code>face</code>
<code>empty_simplex(t)</code>	<code>simplex</code>
<code>face_op(t, s, i)</code>	<code>face</code>
<code>face_op(t, f, i)</code>	<code>edge</code>
<code>face_op(t, e, i)</code>	<code>vertex</code>
<code>vertices(t)</code>	<code>pair<vi, vi></code>
<code>edges(t)</code>	<code>pair<ei, ei></code>
<code>faces(t)</code>	<code>pair<fi, fi></code>
<code>simplices(t)</code>	<code>pair<si, si></code>

Table 3: Requirements of an AS3C. Some remarks about the notation: `t` is a object which type models an AS3C; `empty_vertex(t)` returns a null vertex descriptor; `vi` is the type of a iterator which traverses the vertex container; `vertices(t)` return a pair of iterators: the first points to the first vertex and the second is a “past-the-end” iterator. The other operators work analogously.

As we mentioned previously, we are going to apply the definitions on the requirement tables to simplify the program’s code. Programs 1 and 2 show two versions of the same algorithm, with and without this simplification. This simplified notation will be used in all subsequent programs in the paper.

Program 1 *i*-th vertex of a simplex. Note how the trait class `as3c_traits` isolates the algorithm from the data structure implementation.

```
template <typename T>
as3c_traits<T>::vertex_descriptor ith_vertex(T t, as3c_traits<T>::simplex_descriptor s, int i) {
    int table[][3]={ 1,1,1, 1,1,0, 1,0,0, 0,0,0 };
    return face_op(t, face_op(t, face_op(t, s, table[i][2]), table[i][1]), table[i][0]);
}
```

Program 2 Simplified version of program 1

```
vertex ith_vertex(T t, simplex s, int i) {
    int table[][3]={ 1,1,1, 1,1,0, 1,0,0, 0,0,0 };
    return face_op(t, face_op(t, face_op(t, s, table[i][2]), table[i][1]), table[i][0]);
}
```

Program 2 shows how to get the *i*-th vertex of a simplex and program 3, the *i*-th edge (with respect to the lexicographic ordering over its vertices). The correctness can be verified by looking to figure 1. Of course, we can specialize both programs if more information about the underlying data structure is known.

Program 3 *i*-th edge of a simplex.

```
edge ith_edge(T t, simplex s, int i) {
    int table[][2]={2,2, 2,1, 1,1, 2,0, 1,0, 0,0};
    return face_op(t, face_op(t, s, table[i][1]), table[i][0]);
}
```

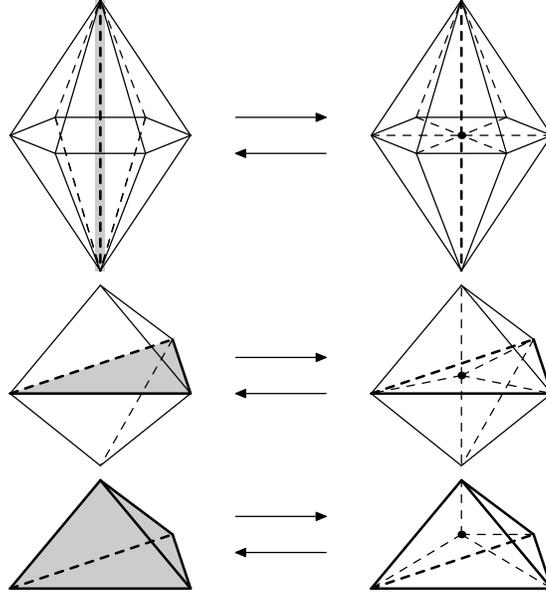


Figure 2: Stellar moves applied to edge, face and simplex, from top to bottom. (\rightarrow) indicates subdivision and (\leftarrow) indicates welding.

3.2 Abstract 3-Manifold

In principle, the requirements on an AS3C are enough to answer incidence queries like “get all faces meeting an edge” or “get all edges meeting a vertex”. But, in many cases, we have more information about the local structure of the complex in each vertex. That information can be used to speed-up those queries. To describe precisely that local structure, we need some definitions.

Two simplices σ_1, σ_2 are *independent* if $\sigma_1 \cap \sigma_2 = \emptyset$. The *join* $\sigma_1 \star \sigma_2$ of independent simplices σ_1, σ_2 is the set $\sigma_1 \cup \sigma_2$. The join of complexes K and L , written $K \star L$, is $\{\sigma \star \tau : \sigma \in K, \tau \in L\}$. The *link* of simplex $\sigma \in K$, denoted $\text{link}(\sigma, K)$, is defined by

$$\text{link}(\sigma, K) = \{\tau \in K : \sigma \star \tau \in K\}.$$

And finally, the *star* of σ in K , $\text{star}(\sigma, K)$, is the join $\sigma \star \text{link}(\sigma, K)$.

The link and star operators provides a combinatorial description of a neighborhood of a simplex. We can use them also to define certain changes in a complex, but care must be taken to not modify essentially (“topologically”) that neighborhood.

The *stellar moves* are a such change. Indeed, many concepts of combinatorial topology are founded on stellar moves [17]. Let K be a complex on the vertex set V , K' a complex on V' , σ a simplex in K and ν a vertex in V' . The operation that changes K into K' by removing $\text{star}(\sigma, K)$ and replacing it with $\nu \star \partial\sigma \star \text{link}(\sigma, K)$ is called a *stellar subdivision* and is written $K' = (\sigma, \nu)K$. The inverse operation $(\sigma, \nu)^{-1}$ that changes K' into K is called a *stellar weld*. These operations are depicted in figure 2.

Two complexes are *stellar equivalent* if they are related by a sequence of stellar moves. A (abstract) *n-ball* is a complex stellar equivalent to a n -simplex and a (abstract) *n-sphere* is a complex stellar equivalent to the boundary of a $(n + 1)$ -simplex.

We can now define a special kind of abstract simplicial complex that has nice local properties.

Definition 2. An abstract n -manifold M is an abstract simplicial n -complex such that for each vertex $\nu \in M$, $\text{link}(\nu, M)$ is a $(n - 1)$ -ball or a $(n - 1)$ -sphere.

The *boundary* of M , denoted by ∂M , is the subcomplex $\partial M = \{\sigma \in M : \text{link}(\sigma, M) \text{ is a ball}\}$. One can proof that ∂M is a $(n - 1)$ -manifold.

description	type
Incident faces iterator	<code>a3m_traits<T>::radial_face_iterator</code>
Incident simplices iterator	<code>a3m_traits<T>::radial_simplex_iterator</code>

Table 4: Associated types of an A3M.

symbol	definition
<code>rfi</code>	<code>typedef a3m_traits<T>::radial_face_iterator rfi;</code>
<code>rsi</code>	<code>typedef a3m_traits<T>::radial_simplex_iterator rsi;</code>

Table 5: A3M related notation.

Many properties follows from definition 2. In our particular case, we want properties that help us to speed-up local queries over abstract 3-manifolds (A3M). Each boundary face of a 3-manifold, for example, is incident to a unique simplex and internal faces (faces not on boundary) are shared by exactly two simplices. Therefore, we can require operators that, in constant time, retrieve all simplices meeting at a face.

Table 6 lists a set of additional requirements to the AS3M ones that are sufficient to formalize the A3M concept. With that additional requirements, we can implement program 4 which, given an internal face `f` and an incident simplex `s`, returns the simplex that shares `f` with `s`.

Program 4 Opposite simplex. The tie function is just a compact way of assign a pair of values to two variables.

```
simplex opposite_simplex(T t, simplex s, face f) {
    simplex s1, s2;
    tie(s1, s2)=incident_simplices(t, f);
    if(s1==s) return s2 else return s1;
}
```

That modest operation is the key component of a *radial iterator*, that is, an iterator that traverses all faces or simplices meeting an edge. Radial iterators are used in algorithms that compute the star of vertices and edges, for instance, and are reminiscent of the Weiler’s radial edge structure (RED) [32]. Again, nothing prevents users from implementing *ad hoc* iterators, perhaps based on some reliable implementation of the facet-edge structure of Dobkin and Laszlo [7].

3.3 Oriented Abstract 3-Manifold

Orientation is another notion we want capture. Since orientation can be defined in a purely combinatorial way, without reference to geometrical concepts, we choose to place the *oriented abstract 3-manifold* concept as a refinement of abstract 3-manifold.

An *orientation* on a n -manifold M is a function s that assigns for each n -simplex $\sigma \in M$, an integer in the set $\{+1, -1\}$. The choice of orientation in σ induces an orientation in its faces in the following way:

$$s(d_i(\sigma)) = (-1)^i s(\sigma), i = 0, \dots, n. \quad (2)$$

An orientation is *coherent* if *contiguous* n -simplices, i.e., simplices sharing an $(n - 1)$ -simplex, induces opposites orientations in its common face, that is,

$$d_i(\sigma_1) = d_j(\sigma_2) \Rightarrow s(d_i(\sigma_1)) = -s(d_j(\sigma_2)),$$

where σ_1 and σ_2 are n -simplices in M . Now, we can define another basic object.

Refinement of abstract simplicial 3-complex	
expression	return type
on_boundary(t, v)	bool
on_boundary(t, e)	bool
on_boundary(t, f)	bool
incident_simplex(t, f)	simplex
incident_simplices(t, f)	pair<simplex, simplex>
a_incident_face(t, e)	face
boundary_faces(t, e)	pair<face, face>
a_incident_edge(t, v)	edge
radial_simplices(t, e)	pair<rsi, rsi>
radial_faces(t, e)	pair<rfi, rfi>

Table 6: Requirements of an A3M. Type t models an A3M. Types rsi and rfi model a radial simplex iterator and radial face iterator, respectively. Some pre-conditions must hold: `boundary_faces(t, e)` can be used only if `on_boundary(t, e)==true`, for instance.

Refinement of abstract 3-manifold	
expression	return type
simplex_orientation(t, s)	int

Table 7: Requirements of an OA3M.

Definition 3. An oriented abstract n -manifold is an abstract n -manifold plus a coherent orientation.

In the 3-dimensional case, the additional requirement on an A3M is just the operator `simplex_orientation` that takes a manifold and a simplex and returns an `int` in the set $\{-1, 1\}$. Program 5 is the obvious implementation of the equation 2.

Program 5 Induced orientation of the faces.

```
int face_orientation(T t, simplex s, face f) {
  int o=simplex_orientation(t, s);
  for(int i=0; i<4; ++i, o*=-1)
    if(face_op(t, s, i)==f) return o;
}
```

3.4 3-Polyhedron and Combinatorial 3-Manifold

Until now, we discussed only combinatorial concepts. Let's introduce the geometrical counterpart of the previously defined concepts.

We call an *euclidean embedding* of an abstract simplicial complex K a function g from the vertex set V to an euclidean space E^m that maps a vertex $\nu \in V$ to a euclidean point $g(\nu) \in E^m$, such that $g(\sigma)$ is a set in general position in E^m , for all $\sigma \in K$. A subset \mathcal{P} of E^m is a *geometric realization* of K if there is an embedding g satisfying

$$x \in \mathcal{P} \Leftrightarrow x \in \text{ConvHull}(g(\sigma)), \text{ for some } \sigma \in K.$$

Below, we define the geometrical objects corresponding to abstract simplicial n -complex and abstract n -manifold.

Definition 4. A n -polyhedron is a set $\mathcal{P} \subset E^m$ for which exists an abstract simplicial n -complex K and an euclidean embedding g such that $\mathcal{P} = |K|_g$.

description	type
Point type	poly3_traits<T>::point_type

Table 8: Associated types of a Poly3.

symbol	definition
point	typedef poly3_traits<T>::point_type point;

Table 9: Poly3 related notation.

Definition 5. A combinatorial n -manifold is a set $\mathcal{M} \subset E^m$ for which exists an abstract n -manifold M and an euclidean embedding g such that $\mathcal{M} = |M|_g$.

From the computational side, a polyhedron concept (Poly3) is just a refinement of AS3C with an additional requirement `euclidean_point` that takes a manifold and a vertex and return an euclidean point, see table 10. A combinatorial 3-manifold (C3M) is a Poly3 plus the requirements of an A3M and an oriented combinatorial 3-manifold (OC3M) is a Poly3 plus the requirements of an OA3M.

Refinement of abstract simplicial 3-complex	
expression	return type
<code>euclidean_point(t, v)</code>	point

Table 10: Requirements of a Poly3.

3.5 Binary Multitriangulation

Now, we'll investigate the interplay between combinatorial and geometrical concepts related to subdivision process and how this leads naturally to the concept of binary multitriangulations.

A polyhedron $\mathcal{P}' = |K'|_g$ is a *subdivision* of the polyhedron $\mathcal{P} = |K|_h$, denoted by $\mathcal{P}' < \mathcal{P}$, if $\mathcal{P}' = \mathcal{P}$ and for each $\sigma' \in K'$ exists a $\sigma \in K$ such that

$$\text{ConvHull}(g(\sigma')) \subset \text{ConvHull}(h(\sigma)).$$

The above definition uses geometrical concepts like euclidean embeddings. Therefore, we can not assert *a priori* anything about how the complexes K and K' are related. However, a theorem of Newman, presented in modern form in [17], shows that $\mathcal{P}' < \mathcal{P}$ if, and only if, K' is stellar equivalent to K . Moreover, the stellar equivalence can be choosed in such a way that only stellar moves on 1-simplices (“edges”) are used.

There is a good reason to restrict the stellar moves to moves on edges. Whenever a stellar subdivision happens in an edge ε , all simplices containing ε are splitted in two. Accordingly, a sequence of stellar subdivision induces a binary tree structure in the simplices. And binary trees often leads to simpler algorithms.

In order to define the binary multitriangulation concept (BMT), we need some auxiliary definitions. We follow closely the definitions in [12]. A *partially ordered set* (poset) $(C, <)$ is a set C with an antisymmetric and transitive relation $<$ defined on its elements. Given $c, c' \in C$, notation $c < c'$ means $c < c'$ and there in no $c'' \in C$ such that $c < c'' < c'$. An element $c \in C$, such that for all $c' \in C$, $c \leq c'$, is called a *minimal* element in C . If there is a unique minimal element $c \in C$, then c is called the *minimum* of C . Analogously are defined *maximal* and *maximum* elements.

Definition 6. A binary multitriangulation is a poset $(\mathcal{T}, <)$, where \mathcal{T} is a finite set of abstract 3-manifolds (named triangulations) and the order $<$ satisfies:

1. $M' < M$ if, and only if, $M' = (\varepsilon, \nu)M$, for some edge $\varepsilon \in M$.
2. There is maximum and minimum abstract 3-manifolds in \mathcal{T} , called base triangulation and full triangulation, respectively ²;

Property 2 says, in fact, that a BMT is a *lattice*. Other fact which follows from the definition is that every two triangulations in \mathcal{T} are stellar equivalent. As usual, a BMT can be thought as a directed acyclic graph (DAG), with one drain and one source, whose arrows are labeled with stellar subdivisions on edges. From an algorithmic perspective, the key idea is to use the above mentioned binary tree structure in the simplices to encode the DAG.

To describe the requirements on a BMT, we need to do a little digression about *state changes* in a data structure. The formerly defined requirements, like `incident_simplices`, are deterministic functions without side effects, at least from the user viewpoint. In other words, `incident_simplices` must return the same value in successive invocations. The situation changes in the BMT case, because we want to be able to move from a triangulation in \mathcal{T} to another. We can regard this move as a state change in the underlying data structure modeling a BMT. The point is that, between state changes, the functions like `incident_simplices` behave deterministically.

The BMT requirements in table 11 are divided in two groups: operators that changes the state (`subdivide`, `weld` and `base_triangulation`) and the others. The operator `base_triangulation` set the current triangulation in \mathcal{T} to the base triangulation, while `subdivide(t, e)` applies a stellar subdivision to the edge `e` and `weld(t, v)` applies a stellar weld “removing” the vertex `v`. The predicate `is_current` is useful to check if a simplex belongs to the current triangulation.

We must remark that operators `subdivide` and `weld` implements just “local” transitions in the DAG, that is, if \mathcal{T} and \mathcal{T}' are the triangulations before and after `subdivide` be called, respectively, then $\mathcal{T}' < \mathcal{T}$. Program 6 illustrates how `subdivide` can be used to achieve non-local transitions. Note that the order of subdivision of the incident simplices to the subdivision edge is relevant: they are subdivided from the lowest to the highest level, where the `level` function is defined in program 7.

Program 6 Non-local subdivide. This program also illustrates how the encoding of traversal capabilities in radial iterators makes the algorithm more generic and readable at no extra cost.

```
void non_local_subdivide(T t, edge e) {
    typedef set<pair<int, edge> > edge_set;
    edge_set sub_edges;
    if((e==empty_edge(t)) || (!was_subdivided(t, e))) return;
    do {
        sub_edges.clear();
        rsi i, end;
        for(tie(i, end)=radial_simplices(t, e); i!=end; ++i) {
            simplex cs=*i;
            edge se=subdivided_edge(t, cs);
            if(se!=e) sub_edges.insert(make_pair(level(t, cs), se));
        }
        for(edge_set::iterator j=sub_edges.begin(); j!=sub_edges.end(); ++j)
            non_local_subdivide(t, j->second);
    } while(!sub_edges.empty());
    subdivide(t, e);
}
```

²One can replace $M' < M$ by $M < M'$ in property 1. In this case, we must interchange base triangulation and full triangulation. This is a transformation from an *increasing* to a *decreasing* BMT. In [22], Puppo demonstrates that increasing and decreasing multitriangulations are equivalent.

Program 7 The simplex level.

```
int level(T t, simplex s) {  
    if(!has_parent(t, s)) return 0;  
    else return level(t, parent(t, s))+1;  
}
```

Refinement of abstract 3-manifold	
expression	return type
was_subdivided(t, e)	bool
subdivide(t, e)	void
in_base_triangulation(t, v)	bool
weld(t, v)	void
base_triangulation(t)	void
has_children(t, s)	bool
children(t, s)	pair<simplex, simplex>
has_parent(t, s)	bool
parent(t, s)	simplex
subdivided_edge(t, s)	edge
welded_vertex(t, s)	vertex
is_current(t, v)	bool
is_current(t, e)	bool
is_current(t, f)	bool
is_current(t, s)	bool

Table 11: Requirements of a BMT. Type t models a BMT. The binary tree structure in the simplices can be traversed with `children` and `parent`.

4 Models

In this section, we present data structures which are models of the concepts defined above, in the sense that they fill all necessary requirements. We have absolutely no pretension of describing “the best” data structure, for the reason that we think that data structures are somehow application dependent. But, if we did a good analysis in the previous section, most algorithms can be used in different applications without change.

Figure 3 resumes the prototypical data structure which models a BMT. Most operators are easily inferred by inspection. Stripping out some data, we obtain models to simpler concepts like AS3C and A3M. It remains to clarify certain points:

- Each element has sufficient information to recover its incident elements (`incident_*` fields) and its star (`star` fields);
- The subdivided edge of a simplex s is given by `ith_edge(t, s, s.subdivided_edge)`;
- The array `edge.star` stores the boundary faces incident to a boundary edge, or stores a face incident to an internal edge;
- A face f is on boundary if, and only if, `f.star[1]==0`, and an edge e is on boundary if, and only if, `e.star[0]!=e.star[1]`.
- A face f is current if, and only if, `f.star[0].current==true`. And an edge e is current if, and only if, `is_current(t, e.star[0])==true`.
- An ordering is adopted in the vertices of the simplex in such a way that the welded vertex is always the last vertex.

Actually, the data structure exhibited in figure 3 it's not quite the same we use in our implementation. The main difference is that the `struct bmt` is parametrized by an *attribute class*, which enables that new attributes be added to any element. One can, for instance, add the plane equation to each face or a scalar value to each vertex. This is done in compile time and causes no storage overhead. A similar technique is described in [3].

The creation process of binary multitriangulations in our implementation is straightforward. There is a `add_simplex` function that takes a triangulation `t` and four vertices `v0`, `v1`, `v2` and `v3` as arguments and adds to `t` all faces and edges not yet in `t`, updating the incidence and star data. Thus, all simplices of the base triangulation are entered. Afterward, according to applications needs, a sequence of calls to `edge_split` is dispatched. The `edge_split` function takes a triangulation `t`, an edge `e` and a *visitor object* `vis` as arguments. It applies a stellar move on `t` that splits the edge `e`. Each time a new element is created, a internal face of a incident simplex for example, a corresponding function in `vis` is called. This allows that user attributes, as the above mentioned face plane equation, be updated. The Visitor concept is explained in [26].

We note that this data structure is too general. Once more, the application guides the real implementation. In dealing with regular spatial decompositions, for example, most of work can be done procedurally (see, for example, [16]). Even out-of-core techniques (e.g., [8, 11]) can be implemented without changing the interface.

```

struct vertex;
struct edge;
struct face;
struct simplex;
typedef vertex * vertex_descriptor;
typedef edge * edge_descriptor;
typedef face * face_descriptor;
typedef simplex * simplex_descriptor;

struct vertex {
    edge_descriptor star;
    bool current;
    bool boundary;
    bool in_base_triangulation;
};

struct edge {
    vertex_descriptor incident_vertices[2];
    face_descriptor star[2];
    bool was_subdivided;
};

struct face {
    edge_descriptor incident_edges[3];
    simplex_descriptor star[2];
};

struct simplex {
    face_descriptor incident_faces[4];
    simplex_descriptor child[2];
    simplex_descriptor parent;
    int subdivided_edge;
    bool current;
};

struct bmt {
    list<simplex_descriptor> simplex_list;
    list<face_descriptor> face_list;
    list<edge_descriptor> edge_list;
    list<vertex_descriptor> vertex_list;
};

```

Figure 3: Modeling a BMT.

5 Applications

The development of multi-resolution techniques for volume visualization is one of our driving applications. Although multi-resolution techniques for 3D surfaces is a well-developed research area [13], the same is not true for multi-resolution techniques for 3D volumes. This is particularly true for multi-resolution techniques for unstructured volumetric grids (see [5] for a recent survey). In fact, compared to the surface case, it is possible to argue that multi-resolution work in volumetric grids is still at its infancy.

The data structures presented in this paper are aimed at providing a more formal and disciplined way of handling unstructured volumetric grids, which we hope will aid the further development of this area. To prove the usefulness of our concepts, we have implemented two simple applications: a simple unstructured grids volume renderer based on the ideas presented in [33, 25]; and a progressive volume approximation system similar in some respects to the one described by Roxborough and Nielson [24].

Direct Volume Rendering. An efficient technique for exploring graphics hardware for volume rendering is the Projected Tetrahedra (PT) algorithm of Shirley and Tuchman [25], which uses the traditional 3D polygon-rendering pipeline. This technique renders a volumetric grid by breaking the volumetric grid into a collection of tetrahedra. Then, each tetrahedra is rendered by *splatt*ing its faces on the screen. A key idea of the PT algorithm is exploit the fact that the aspect graph [21] of a tetrahedron consists of a few simple cases which can be encoded in a small table which depends completely on the dot products of the normal of the faces of the tetrahedron with the viewing direction. On average, one needs to render 3.4 triangles per tetrahedron [34].

In order to apply PT, one needs to compute a visibility-ordering of the cells. Williams' Meshed Polyhedra Visibility Ordering (MPVO) algorithm [33] developed in the early 1990s provides a very fast visibility-ordering algorithm. The MPVO algorithm explores the topological adjacencies of a convex model to produce a visibility ordering of the cells with respect to a given viewpoint. Given a tetrahedral mesh S , it produces an adjacency graph G of the mesh S , with vertices v of the graph corresponding to each tetrahedral cell c of the mesh, and edges e representing a common face between cells. A directed version of this graph is obtained when a viewpoint or viewing direction is specified, which is then used to define the direction of each edge by a simple dot product with the face normals (see Figure 4). For every new viewpoint, a directed graph as described is formed, and a topological sorting procedure outputs cells in visibility ordering. MPVO, which runs in linear time, works well for well-behaved meshes (acyclic and convex). (General acyclic meshes can be handled with more complex algorithms as shown in [28].)

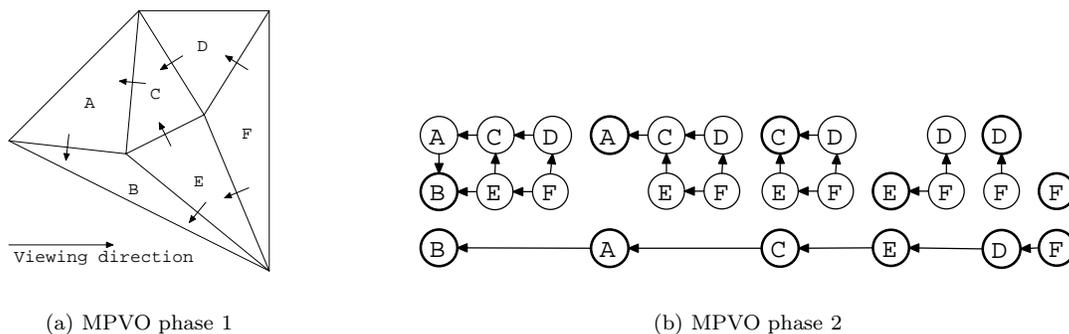


Figure 4: Depiction of the MPVO algorithm [33]. MPVO works by first determining pairwise ordering relations between cells that share a face (a). Then, in a second phase, the ordering is determined by a topological sort of the induced visibility graph (b).

Our implementation of the PT part of the algorithm is quite simple, and it is completely driven by a set of tables for classifying the different cases (including degenerate cases). Figure 5 shows the six cases of projected tetrahedra.

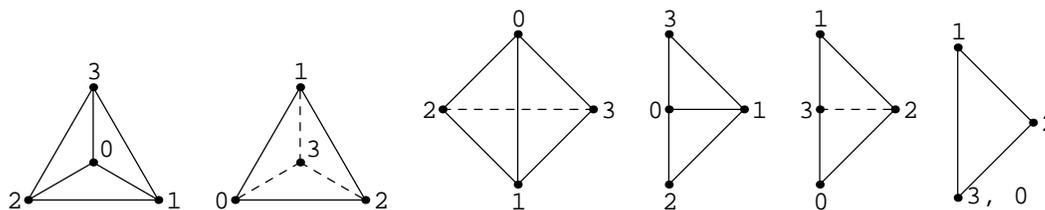


Figure 5: Cases of the Projected Tetrahedra (PT) algorithm [25]. Depending on the particular case, between one to four triangles are rendered.

Implementing the sorting is more interesting, since we are able to completely avoid building a separate graph or actually writing any sorting code. Our implementation for computing the visibility orders works by developing a *visibility graph* concept that is an adaptor over the C3M concept, i.e., it provides the appropriate graph interface to a combinatorial manifold. So, we can use all graph machinery built in the Boost Graph Library [26], for example, to solve the visibility problem. More specifically, in the simpler regular convex case, the visibility is computed by a single call to the BGL function `topological_sort`.

Figure 6 gives an example of direct volume rendering using the algorithm described above. Figure 6 (a) shows the underlying space decomposition and the tetrahedral mesh. Figure 6 (b) shows the volume visualization of a lung tumor dataset.

Progressive Tetrahedral Approximation of 3D Functions. A more interesting example of the use of our framework is our system for the hierarchical approximation and rendering of 3D scalar functions (see [20] for a survey). We use a simple approximation algorithm which recursively refines the triangulation until a user-defined error bound has been achieved. The exact procedure we implemented works by generating a collection of random points inside a given tetrahedron [23]. Then finding the best least-squares linear fit for the scalar function inside the simplex. If the error bound of the approximation is above the user-defined threshold, we continue to subdivide.

The approximation algorithm we implemented is somewhat similar to the described in [24] for computing a progressive tetrahedral approximation of ultrasound data, and it could potentially be used for the same purpose. In their system, a global least-squares fit is used, while in our system we handle each tetrahedron separately. Note that for a given vertex of the triangulation, each incident simplex determines one (possibly different) scalar value. We set the value at the vertex as the average value. Alternatively, we could perform an analysis of the discontinuities in the scalar function and represent this information as radial wedges [15].

For rendering purposes, we use a very simple sorting procedure shown in Program 8. The basic idea is the same used for traversing a BSP in back-to-front order. We always traverse the space region that does not contain the viewpoint (i.e., the “back”), before traversing the region that contains the viewpoint. This leads to a very efficient rendering algorithm.

Program 8 Hierarchical splatting.

```

void hier_splat(T t, simplex s) {
    if(is_current(t, s))
        splat(t, s);
    else {
        face f=internal_face(t, s);
        simplex s0, s1;
        tie(s0, s1)=children(t, s);
        if(visibility_test(t, f)) {
            hier_splat(t, s1);
            hier_splat(t, s0);
        } else {
            hier_splat(t, s0);
            hier_splat(t, s1);
        }
    }
}

```

Figure 7 demonstrates the BMT adaptation capabilities. These adapted decompositions were computed using least square fitting described above. Figure 7 (a) shows the mesh corresponding to a function that varies linearly in one direction and has a sharp discontinuity. Figure 7 (b) shows the mesh corresponding to the characteristic function of a sphere. Observe that we compute the full hierarchical decomposition, where the top level is a subdivision of a cube into six tetrahedra. In these images, we have removed one of the top level tetrahedra to reveal the internal structure of the mesh.

6 Conclusion

We have presented a generic programming framework for multiresolution spatial decompositions which was formulated through a rigorous mathematical analysis of the concepts involved. Indeed, our current implementation contains numerous generic algorithms for the extraction of topological information, as well non-generic functions to build a multiresolution mesh and execute other operations such as input/output. It is certainly possible to employ generic programming techniques in the creation of multiresolution meshes. Nonetheless, the problem is more involved and we plan to consider it in a future paper.

Concerning to future works, there are essentially two lines to follow. The first line is theoretical and consists in to investigate the BMT properties, its n -dimensional extension, its applicability in non-manifold settings and the incorporation of all stellar moves, not only moves on edges. The second line is to proof, by means of applications, mainly in the context of visualization of large datasets and finite element analysis, the advantages of our approach.

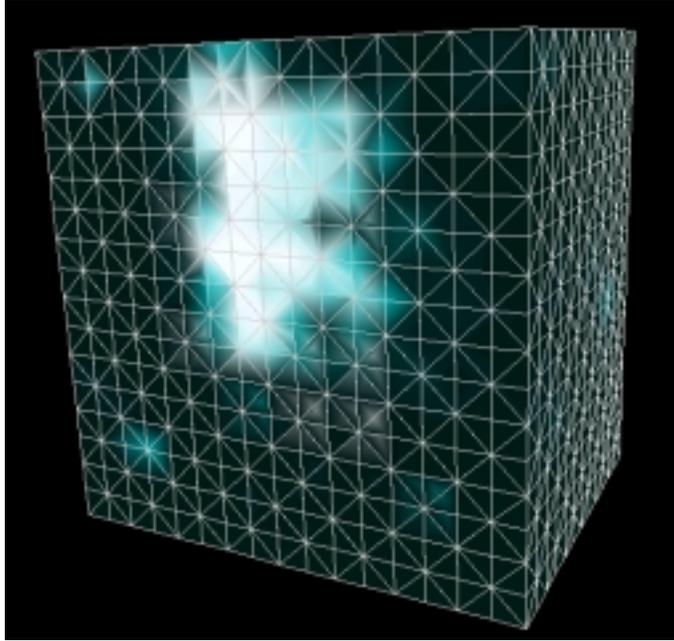
We think that our work can be classified in the confluence of two new trends in graphics. On one hand, our concern to clearly separate geometric and topological concepts, emphasizing the later ones, brings us closer to *Computational Topology* as posed in [6]. This new branch is as promising today as Computational Geometry was thirty years ago. On the other hand, generic programming is a powerful methodology for computer programming, which holds the promise to complete, specially in relation to algorithm abstraction, the revolution started twenty years ago by object-oriented programming. Our hope is that this work will become the basis of a library called CTAL, that is, Computational Topology Algorithm Library.

References

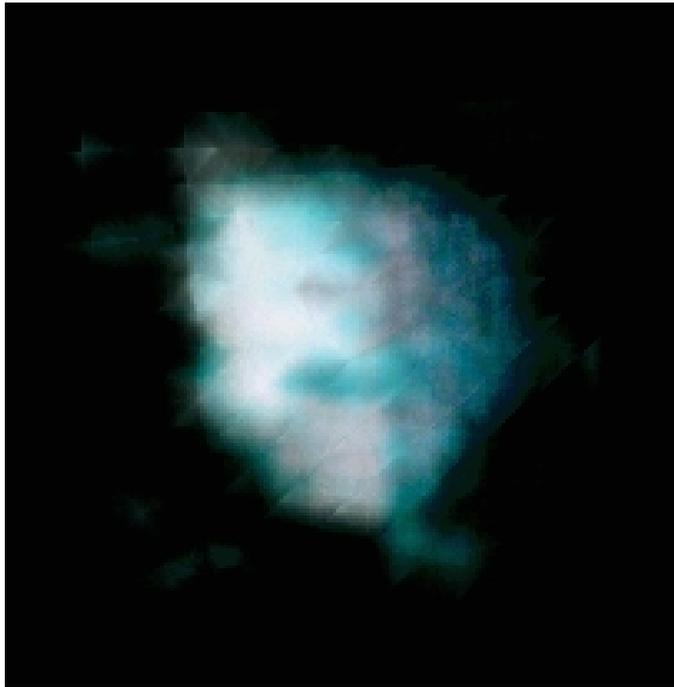
- [1] J. Alexander. The combinatorial theory of complexes. *Ann. Math.*, 31:294–322, 1930.
- [2] G. Berti. *Generic software components for Scientific Computing*. PhD thesis, BTU Cottbus, 2000.
- [3] Mario Botsch, Stephan Steinberg, Stephan Bischoff, and Leif Kobbelt. Openmesh - a generic and efficient polygon mesh data structure. In *OpenSG PLUS Symposium*, 2002.
- [4] P. Cignoni, L. De Floriani, P. Magillo, E. Puppo, and R. Scopigno. TAn2 - visualization of large irregular volume datasets. Technical Report DISI-TR-00-07, University of Genova (Italy), 2000.
- [5] E. Danovaro, L. De Floriani, M. Lee, and H. Samet. Multiresolution tetrahedral meshes: an analysis and a comparison. In *Proceedings International Conference on Shape Modeling*, 2002.
- [6] T. Dey, H. Edelsbrunner, and S. Guha. Computational topology. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry (Contemporary mathematics 223)*, pages 109–143. American Mathematical Society, 1999.
- [7] D. P. Dobkin and M. J. Laszlo. Primitives for the manipulation of threedimensional subdivisions. *Algorithmica*, 4:3–32, 1989.
- [8] Jihad El-Sana and Yi-Jen Chiang. External memory view-dependent simplification. *Computer Graphics Forum*, 19(3), August 2000. ISSN 1067-7055.
- [9] Jihad El-Sana and Amitabh Varshney. Generalized view-dependent simplification. *Computer Graphics Forum*, 18(3):83–94, September 1999. ISSN 1067-7055.
- [10] Andreas Fabri, Geert-Jan Giezeman, Lutz Kettner, Stefan Schirra, and Sven Schonherr. On the design of CGAL a computational geometry algorithms library. *SP&E*, 30(11):1167–1202, 2000.
- [11] Ricardo Farias and Cláudio T. Silva. Out-of-core rendering of large, unstructured grids. *IEEE Computer Graphics & Applications*, 21(4):42–51, July / August 2001.

- [12] L. De Floriani, E. Puppo, and P. Magillo. A formal approach to multiresolution modeling. In W. Straßer, R. Klein, and R. Rau, editors, *Theory and Practice of Geometric Modeling*. Springer-Verlag, 1996.
- [13] M. Garland. Multiresolution modeling: Survey & future opportunities. In *Eurographics '99, State of the Art Report (STAR)*, 1999.
- [14] L. J. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of voronoi diagrams. *ACM Trans. Graph.*, 4:74–123, 1985.
- [15] Hugues Hoppe. Efficient implementation of progressive meshes. *Computers & Graphics*, 22(1):27–36, February 1998. ISSN 0097-8493.
- [16] M. Lee, L. De Floriani, and H. Samet. Constant time neighbor finding in hierarchical tetrahedral meshes. In *Proceedings International Conference on Shape Modeling*, pages 286–295, 2001.
- [17] W. B. R. Lickorish. Simplicial moves on complexes and manifolds. In *Proceedings of the Kirbyfest*, volume 2, pages 299–320, 1999.
- [18] M. Mäntylä. *An Introduction to Solid Modeling*. Computer Science Press, Rockville, Maryland, 1988.
- [19] J. Peter May. *Simplicial Objects in Algebraic Topology*, volume 11. D. Van Nostrand Company, Inc., Princeton, 1967.
- [20] Greg Nielson. Tools for triangulations and tetrahedrizations and constructing functions defined over them. In *Scientific Visualization: Overviews, Methodologies, and Techniques*, pages 419–515. IEEE CS Press, 1997.
- [21] H. Plantinga and C. R. Dyer. Visibility, occlusion, and the aspect graph. *Internat. J. Comput. Vision*, 5(2):137–160, 1990.
- [22] E. Puppo. Variable resolution triangulations. *Computational Geometry Theory and Applications*, 11(34):219–238, 1998.
- [23] C. Rocchini and P. Cignoni. Generating random points in a tetrahedron. *Journal of Graphics Tools*, 5(4):9–12, 2000.
- [24] T. Roxborough and Gregory M. Nielson. Tetrahedron based, least squares, progressive volume models with application to freehand ultrasound data. In *IEEE Visualization 2000*, pages 93–100, October 2000.
- [25] P. Shirley and A. A. Tuchman. Polygonal approximation to direct scalar volume rendering. *Computer Graphics*, 24(5):63–70, 1990.
- [26] Jeremy G. Siek, Lie-Quan Lee, and Andrew Lumsdaine. *Boost Graph Library, The: User Guide and Reference Manual*. Addison-Wesley, 2002.
- [27] Jeremy G. Siek and Andrew Lumsdaine. The matrix template library: A generic programming approach to high performance numerical linear algebra. In *ISCOPE*, pages 59–70, 1998.
- [28] C. T. Silva, J. S. B. Mitchell, and P. Williams. An exact interactive time visibility ordering algorithm for polyhedral cell complexes. In *1998 Volume Visualization Symposium*, pages 87–94, October 1998.
- [29] A. A. Stepanov and M. Lee. The Standard Template Library. Technical Report X3J16/94-0095, WG21/N0482, ISO Programming Language C++ Project, 1994.
- [30] B. Stroustrup. *The C++ Programming Language: Third Edition*. Addison-Wesley, 1997.

- [31] Luiz Velho and Jonas Gomes. Variable resolution 4-k meshes: Concepts and applications. *Computer Graphics forum*, 19:195–212, 2000.
- [32] Kevin Weiler. Edge-Based Data Structures for Solid Modeling in Curved-Surface Environments. *IEEE Computer Graphics and Applications*, 5(1):21–40, 1985.
- [33] P. Williams. Visibility ordering meshed polyhedra. *ACM Transactions on Graphics*, 11(2), 1992.
- [34] C. Wittenbrink. Cellfast: Interactive unstructured volume rendering. In *Proceedings IEEE Visualization'99, Late Breaking Hot Topics*, pages 21–24, 1999. Also available as Technical Report, HPL-1999-81R1, Hewlett-Packard Laboratories.

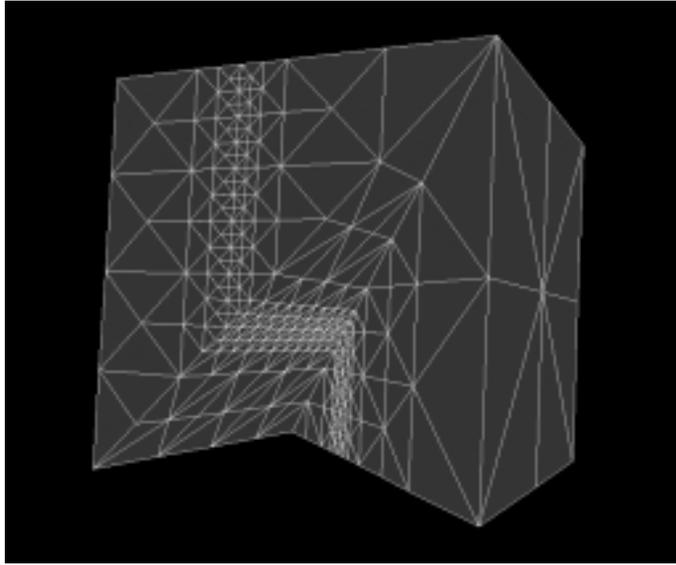


(a)

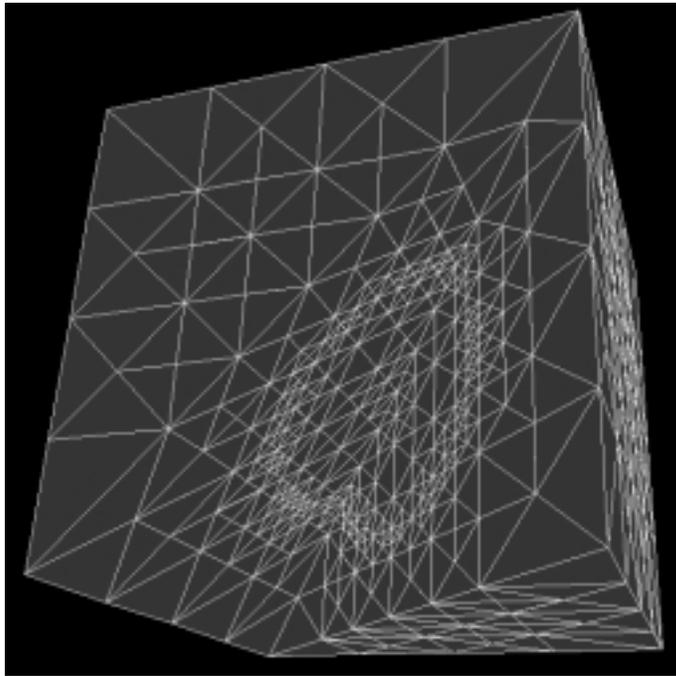


(b)

Figure 6: Volume visualization of tumor dataset.



(a)



(b)

Figure 7: Adapted decompositions of two different scalar functions.