The Implementation of the Colored Abstract Simplicial Complex and Its Application to Mesh Generation

CHRISTOPHER T. LEE, JOHN B. MOODY, ROMMIE E. AMARO,
J. ANDREW MCCAMMON, and MICHAEL J. HOLST, University of California San Diego, CA

We introduce the Colored Abstract Simplicial Complex library (CASC): a new, modern, and header-only C++ library that provides a data structure to represent arbitrary dimension abstract simplicial complexes (ASC) with user-defined classes stored directly on the simplices at each dimension. This is accomplished by using the latest C++ language features including variadic template parameters introduced in C++11 and automatic function return type deduction from C++14. Effectively, CASC decouples the representation of the topology from the interactions of user data. We present the innovations and design principles of the data structure and related algorithms. This includes a metadata-aware decimation algorithm, which is general for collapsing simplices of any dimension. We also present an example application of this library to represent an orientable surface mesh.

CCS Concepts: • **Mathematics of computing** → **Mathematical software**; *Combinatorics*; *Mesh generation*; • **Theory of computation** → **Data structures design and analysis**; **Computational geometry**; • **Applied computing** → **Imaging**; *Molecular structural biology*; • **Computing methodologies** → Computer graphics;

Additional Key Words and Phrases: Abstract simplicial complexes, molecular modeling, mesh generation, mesh decimation, variadic templates, C++ library

ACM Reference format:

Christopher T. Lee, John B. Moody, Rommie E. Amaro, J. Andrew McCammon, and Michael J. Holst. 2019. The Implementation of the Colored Abstract Simplicial Complex and Its Application to Mesh Generation. *ACM Trans. Math. Softw.* 45, 3, Article 28 (August 2019), 20 pages.

<https://doi.org/10.1145/3321515>

C. T. Lee and J. B. Moody contributed equally to this work.

This work is supported by the National Institutes of Health, NIGMS under grant numbers P41-GM103426 and RO1-GM31749. CTL also acknowledges support from the NIH Molecular Biophysics Training Grant T32-GM008326. MJH was supported in part by the National Science Foundation, Division of Mathematical Sciences under awards DMS-CM1620366 and DMS-FRG1262982.

Authors' addresses: C. T. Lee, R. E. Amaro, and J. A. McCammon, Department of Chemistry and Biochemistry, University of California San Diego, La Jolla, CA, 92093 USA; emails: {ctlee, ramaro, jmccammon}@ucsd.edu; J. B. Moody and M. J. Holst, Department of Mathematics, University of California, San Diego, La Jolla, CA, 92093 USA; emails: brogan@gmail.com, mholst@ucsd.edu.

Author's current address: J. B. Moody, ViaSat, Inc., Carlsbad-Bldg 10-2063, 6155 El Camino Real, Carlsbad, CA 92009.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

0098-3500/2019/08-ART28 \$15.00

<https://doi.org/10.1145/3321515>

1 INTRODUCTION

For problems in computational topology and geometry, it is often beneficial to use simple building blocks to represent complicated shapes. A popular block is the simplex, or the generalization of a triangle in any dimension. Due to the ease of manipulation and the coplanar property of triangles, triangulations have become commonplace in fields such as geometric modeling and visualization as well as topological analysis. Discretizations are also used for efficient solving of Partial Differential Equations (PDE). The use of meshes has become increasingly popular even in the fields of computational biology and medicine [24].

As methods in structural biology improve and new datasets become available, there is interest in integrating experimental and structural data to build new predictive computer models [19]. A key barrier that modelers face is the generation of multi-scale, computable, geometric models from noisy datasets such as those from Electron Tomography (ET)[23]. This is typically achieved in at least two steps: (1) segmentation of relevant features, and (2) approximation of the geometry using meshes. Subsequently, numerical techniques such as Finite Elements Modeling or Monte Carlo can be used to investigate the transport and localization of molecules of interest.

While many have studied mesh generation in the fields of engineering and animation, few methods are suitable for biological datasets. This is largely due to noise introduced by limits in image resolution or contrast. Even while using state-of-the-art segmentation algorithms for ET datasets, there are often unresolved or missed features. Due to these issues, the generated meshes often have holes and other non-manifolds that must be resolved prior to mathematical modeling. Another challenge is the interpretation of a voxel valued segmentation. The conversion of zig-zag boundaries into a mesh can lead to other problems such as extremely high aspect ratio triangles or, in general, poorly conditioned elements [23]. To remedy this, various smoothing and decimation algorithms must also be applied prior to simulation.

Previous work by us and others have introduced a meshing tool for biological models, GAMer, for building 3D tetrahedral meshes that obey internal and external constraints, such as matching embedding and/or enclosing molecular surfaces. It also provides the ability to use various mesh improvement algorithms for volume and surface meshes [13, 22]. GAMer uses the Tetgen library as the primary tetrahedral volume generator [20]. While the algorithms are sound, the specific implementation is prone to segmentation faults even for simple meshes. Careful analysis of the code has identified that the data structures used for the representation of the mesh is primarily at fault. This article will focus entirely on the representation of topology in very complex mesh generation codes. We note that the algorithms that handle geometric issues like shape regularity and local adaptivity are well understood [2, 16], among others. Similarly there is a large body of literature related to local mesh refinement and decimation [3, 4]. Our innovations serve to enable the implementation of these algorithms in the most general and robust way.

GAMer currently employs a neighbor list data structure that tracks the adjacency and orientation of simplices. Neighbor lists are quick to construct; however, the representation of non-manifolds often leads to code instability. Algorithms must check for aberrant cases creating substantial overhead. We note that while the need to gracefully represent 2D and 3D non-manifolds for ET applications drove our initial focus, we are also interested in mesh generation in higher dimensions with applications to: numerical general relativity (3D+1) [14, 18], computational geometric analysis (nD) [15], phase space simulations (6D), and arbitrary collective variable spaces in molecular modeling for enhanced sampling [21]. We therefore chose the following requirements for a mesh data structure to serve as design goals:

- General and capable of representing non-manifold, mixed dimensional, oriented, and non-oriented meshes in arbitrary dimensions.

- Support for inline and flexible data storage. In some applications, data must be associated with the topology. For example, problems in general relativity typically require the storage of metric tensors on all simplex dimensions.
- Support for intuitive and simple manipulations and traversals.

Here, we describe the development of a scalable *colored abstract simplicial complex* data structure called CASC. Simplices are stored as nodes on a Hasse diagram. For ease of traversal, all adjacency is stored at the node level. An additional data object can be stored at each node that is typed according to the simplex dimension at compile time. This means that, for example, for a mesh, the 0-simplices can be assigned a vertex type while the 2-simplices can store some material property instead. Typing of each k -simplex is achieved using variadic templates introduced in C++11. CASC thus provides a natural separation between the combinatorics represented by the ASC from the underlying data types at each simplex dimension and their interactions. In Section 2, we briefly define an ASC and some relevant definitions followed by the introduction of the CASC data structure and its construction in Section 3. We then demonstrate the use of CASC to represent a surface mesh and compute vertex tangents in Section 5.

1.1 Related Work

Although many data structures to represent simplicial complexes have been developed, to the best of our knowledge, there currently exists no data structure that supports meshes of arbitrary dimension with user-selected typed data stored directly on each simplex. A full review of all existing data structures is beyond the scope of this work; however, we highlight several representative examples. Many data structures such as the half-edge and doubly-connected edge list among others are restricted to the representation of two-manifolds only [9]. Other data structures such as Simplified Incidence Graph (SIG) [8], Incidence Simplicial (IS) [10], Generalized Indexed Data Structure with Adjacencies (IA*) [7], SimplexTree [6], Array-based Half-Facet (AHF) [11, 17], LinearCellComplex, and dDimensional Triangulations [5] support or can be extended to represent arbitrary dimensional simplicial complexes. However, their current implementations either do not consider the storage of data beyond possibly embedding or do not support inline storage of user data. LinearCellComplex from the Computational Geometry Algorithms Library (CGAL) supports only a linear geometrical embedding [1]. AHF implemented in Mesh Oriented datABase (MOAB) uses separate arrays of data, which are then referenced using a handle [11, 17]. In addition to the limitations of data storage, some make assumptions limiting their generality. dD Triangulations, for example, assume that a simplicial complex is pure and therefore does not support the representation of mixed dimensional complexes [5].

2 BACKGROUND—ABSTRACT SIMPLICIAL COMPLEXES

An Abstract Simplicial Complex (ASC) is a combinatorial structure that can be used to represent the connectivity of a simplicial mesh, independent of any geometric information. More formally, the definition of an ASC is as follows.

Definition 2.1. Given a vertex set V , an *abstract simplicial complex* $\mathcal{F}$ of V is a set of subsets of V with the following property: for every set $X \in \mathcal{F}$, every subset $Y \subset X$ is also a member of $\mathcal{F}$.

The sets $s \in \mathcal{F}$ are called a simplex or face of $\mathcal{F}$; similarly a face X is said to be a face of simplex s if $X \subset s$. Since X is a face of s , s is a coface of X . Each simplex has a dimension characterized by $\dim s = |s| - 1$, where $|s|$ is the cardinality of set s . A simplex of $\dim s = k$ is also called a k -simplex. The dimension of the complex, $\dim(\mathcal{F})$, is defined by the largest dimension of any member face. Simplices of the largest dimension, $\dim(\mathcal{F})$ are referred to as the facets of the complex.

If one simplex is a face of another, they are incident. Every face of a k -simplex s with dimension $(k - 1)$ is called a boundary face while each incident face with dimension $(k + 1)$ is a coboundary

Fig. 1. Hasse diagrams of several Abstract Simplicial Complexes and a geometric realization from left to right: the empty set, a vertex, an edge, a triangle, a tetrahedron.

face. Two k -simplices, f and s are considered adjacent if they share a common boundary face or coboundary face. The boundary of simplex s , ∂s , is the sum of the boundary faces.

Having introduced the concept of an ASC, we can also define several operations useful when dealing with ASCs. A subcomplex is a subset that is a simplicial complex itself. The Closure (Cl) of a simplex, f , or some set of simplices $F \subseteq \mathcal{F}$ is the smallest simplicial subcomplex of $\mathcal{F}$ that contains F :

$$\text{Cl}(f) = \{s \in \mathcal{F} \mid s \subseteq f\}; \quad \text{Cl}(F) = \bigcup_{f \in F} \text{Cl}(f) \quad (\text{closure}). \quad (1)$$

It is often useful to consider the local neighborhood of a simplex. The Star (St) of a simplex f is the set of all simplices that contain f :

$$\text{St}(f) = \{s \in \mathcal{F} \mid f \subseteq s\}; \quad \text{St}(F) = \bigcup_{f \in F} \text{St}(f) \quad (\text{star}). \quad (2)$$

The Link (Lk) of f consists of all faces of simplices in the closed star of f that do not intersect f :

$$\text{Lk}(f) = \{s \in \text{Cl} \circ \text{St}(f) \mid s \cap f = \emptyset\} = \text{Cl} \circ \text{St}(f) - \text{St} \circ \text{Cl}(f) \quad (\text{link}). \quad (3)$$

For some algorithms, it is often useful to iterate over the set of all vertices or edges, and so on. We use the following notation for the horizontal “level” of an abstract simplicial complex.

$$\text{Lvl}_k(\mathcal{F}) = \{s \in \mathcal{F} \mid \dim s = k\} \quad (4)$$

A subcomplex that contains all simplices $s \in \mathcal{F}$ where $\dim(s) \leq k$ is the k -skeleton of $\mathcal{F}$:

$$\mathcal{F}_k = \text{Cl} \circ \text{Lvl}_k(\mathcal{F}) = \bigcup_{i \leq k} \text{Lvl}_i(\mathcal{F}). \quad (5)$$

By Definition 2.1, an ASC forms a partially ordered set, or poset. Posets are frequently represented by a Hasse diagram, a directed acyclic graph, where nodes represent sets, and edges denote set membership. Several example simplicial complexes and their corresponding Hasse diagrams are shown in Figure 1. Colloquially, we will use *up* and *down* to refer to the coboundary and boundary of a simplex, respectively. In Hasse diagrams, we follow a convention that simplices shown

Fig. 2. Data structure diagram of a triangle represented using CASC. Each simplex is represented as a node containing a dictionary up and/or down, which maps the vertex index to a pointer to the next simplex. Data can be stored at each node with type determined at compile time. Effectively, each level can contain different metadata as defined by the user, separating the interactions of user data from the representation of topology.

graphically on the same horizontal level have the same simplex dimension. Furthermore, simplices of greater dimension are drawn above lesser simplices.

3 COLORED ABSTRACT SIMPLICIAL COMPLEX

In this section, we introduce the CASC data structure and its implementation. For a given simplicial complex, each simplex is represented by a node (`asc_Node`) in the Hasse diagram and defined by a set of keys corresponding to the vertices that comprise the simplex. Note that we use node to refer to objects in CASC Hasse diagram and not 0-simplices. Instead, 0-simplex are referred to as the vertices of the mesh. Furthermore, we refer to the $\emptyset$ -simplex or -1 -simplex as the root simplex interchangeably. When a node is instantiated, we assign it a unique Integer Internal Identifier (`iID`) for use in the development of CASC algorithms. The `iID` is constant and never exposed to the end-user except for debugging purposes. Instead, nodes can be referenced by the user using the `SimplexID`, which acts as a convenience wrapper around an `asc_Node*`, providing additional support for move semantics for fast data access. All topological relations (i.e., edges of the Hasse diagram) are stored in each node as a dictionary, which maps user specified keys to `SimplexIDs` up and down. An example data structure diagram of triangle $\{1,2,3\}$ is shown in Figure 2. Based upon this example, if a user has the `SimplexID` of 1-simplex $\{1,2\}$ and wishes to get 2-simplex $\{1,2,3\}$, they can look in the Up dictionary of `SimplexID`{1,2} for key 3, which maps to a `SimplexID`{1,2,3}. The vertices that constitute each simplex are not stored directly but can be accessed by aggregating all keys in Down.

We note that while the representation of all topological relations is redundant and may not be memory optimal, it vastly simplifies the traversals across the complex. Furthermore, the associate algorithms and innovations using variable typing are general and thus compatible with other more condensed representations.

Fig. 3. Template arguments for ASC are unpacked and assigned to nodes accordingly. The first argument “int” is the key type for labeling vertices while the following arguments define node data types. Notably, Node<2> does not allocate memory for data as the corresponding template argument was “void.” By passing in additional types, simplicial complexes of higher dimensions are instantiated.

3.1 Variable Typing Per Simplex Dimension

We achieve coloring by allowing user-defined data to be stored at each node. The typical challenge for strongly typed languages such as C++ is that the types must be defined at compile time. Typical implementations would either hard code the type to be stored at each level or use a runtime generic type, such as `void*`. However, each of these have drawbacks. For the former, this requires writing a new node data structure for every simplicial complex we may wish to represent. For the latter, using `void*` adds an extra pointer dereference, which defeats cache locality and may lead to code instability. Another possible implementation might be to require users of the library to derive their data types from a common class through inheritance. This solution puts an unnecessary burden on users who may have preexisting class libraries or simply wish to store a built in type, such as an `int`. Furthermore, under the inheritance scheme, changes to the underlying container may require users to update their derived classes. To avoid this cumbersome step, we have employed the use of variadic templates introduced in C++11 to allow for unpacking and assignment of data types. The user specifies the types to be stored at each level in a list of templates to the object constructor; see Figure 3.

The variadic templating allows CASC to represent complexes of any user-defined dimension. To specify an N -simplicial complex, CASC requires the definition of an index/key type followed by $N + 1$ data types and N edge types. The first data type provided after the key type corresponds to data stored on the $\emptyset$ -simplex, which can be thought of as global metadata. For example, suppose we have a 2-simplicial complex intended for visualization and wish to store locations of vertices and colors of faces. A suitable ASC can be constructed using the following template command:

```
auto mesh = AbstractSimplicialComplex<int, void, Vertex, void, Color>()
```

If we now wish to represent a tetrahedral mesh, instead of constructing a new data structure, we can simply adjust the command:

```
auto mesh = AbstractSimplicialComplex<int, void, Vertex, void, Color, Density>()
```

In both cases, the first template argument is the key type for referring to vertices followed by the data type for each k -simplex. Supposing that the user does not wish to store data on any given level, by passing “void” as the template argument, the compiler will optimize the node data type and no memory will be allocated to store data. In both cases, the 0- and 1-simplices will have no data.

By using variadic templates, we allow the user to specify both the dimension of the simplicial complex as well as the types stored at each level. Because the type deduction is performed at compile time, there is no runtime performance impact on user codes. There is, however, some additional code complexity introduced. If the user wishes to retrieve the data stored in a simplicial complex, they must know what level they are accessing at compile time. A consequence is that the exposed identifier object, `SimplexID`, is templated on the integral level so that types can be deduced. This does not present a problem for simple use cases, such as:

```
// Create alias for 2-simplicial complex
using ASC = AbstractSimplicialComplex <int,int,int,int>;
ASC mesh = ASC(); // construct the mesh object
mesh.insert<2>({1,2},5); // insert edge {1,2} with data 5
ASC::SimplexID<2>s = mesh.get_simplex<2>({1,2}); // get the edge
// NOTE: the type is templated on the integral level
std::cout << *s << std::endl; // prints data "5"
```

However, when implementing algorithms intended to be generic on any simplicial complex, templated code must be written. We discuss the implementation of several such algorithms in the following section.

4 IMPLEMENTED ALGORITHMS

The following algorithms are provided with the CASC library:

- Basic Operations
 - Creating and deleting simplices (insert/remove)
 - Searching and traversing topological relations (GetSimplexUp/GetSimplexDown)
- Traversals
 - By level (get_level)
 - By adjacency (neighbors_up/neighbors_down)
 - Traversals across multiple node types (Visitor Design Pattern/double dispatch)
- Complex Operations
 - Star/Closure/Link
 - Metadata aware decimation

4.1 Basic Operations

4.1.1 Creating and Deleting Simplices. Since the CASC data structure maintains every simplex in the complex and all topological relations, inserting a k -simplex, s , into the complex means ensuring the existence of, and possibly creating, $O(2^k)$ nodes and $O(k \cdot 2^{k-1})$ edges (see derivation in SI). Fortunately, the combinatorial nature of simplicial complexes allows this to be performed recursively. A generalized recursive insertion operation for any dimensional complex and user specified types is described in Algorithm 1. The insertion algorithm defines an insertion order such that all dependent simplices exist in the complex prior to the insertion of the next simplex.

As an illustrative example of the template code used in this library, Algorithm 1 is rewritten in C++ template function-like pseudocode shown in Algorithm 6. While the templated code is more complicated, it provides many optimizations. For example, since the looping and recursion are performed at compile time, for any k -simplex we wish to insert, any modern compiler should optimize the code into a series of `insertNode()` calls; the `setupForLoop()` and `forLoop()` function calls can be completely eliminated. As a result, the optimized templated code will exhibit superior runtime performance. To illustrate the insertion operation, a graphical representation of inserting

Fig. 4. Recursive insertion of tetrahedral simplex $\{1,2,3,4\}$. The order of node insertions is represented by the numbered red arrows. When each node is created, the black arrows to parent simplices are created by backfilling.

ALGORITHM 1: Insertion of a new simplex

Input: $keys[n]$: Indices of n simplices to describe new simplex s

rootSimplex: The simplex to insert relative to (most commonly *root*)

Output: The new simplex s

Function $insert(keys[n], rootSimplex)$

```

{
  for ( $i = 0; i < n; i++$ )
  {
    newSimplex = createSimplex( $rootSimplex \cup keys[i]$ )
    if ( $i > 0$ ) /* Recurse to insert sub-simplices */
    {
      /* Pass only the first part of  $keys$  up to index  $i$  */
      return insert( $keys[0:i], newSimplex$ )
    }
    else /* Terminal conditional */
    {
      return newSimplex
    }
  }
}

```

tetrahedron $\{1,2,3,4\}$ by Algorithm 6 is shown in Figure 4, and step-by-step in Figure S6. In the example, new simplex $root \cup v$ is added sequentially to the complex and any missing topological relations are found by traversing the faces of *root* and backfilling.

The removal of any simplex is also performed using a recursive template function. When removing simplex s , in order to maintain the property of being a simplicial complex, all cofaces of s

ALGORITHM 2: Searching for a new simplex**Input:** *root*: the simplex to start the search from.*keys*[*n*]: the relative name of the desired simplex *s***Output:** The simplex *s***Function** *getSimplexUp*(*root*, *keys*[*n*], *index*)

```

{
  if (index < n)
  {
    Find keys[index] in root.up /* Look for the edge up */
    if (Edge up is found)
    {
      return getSimplexUp(root ∪ keys[index], keys, index+1)
    }
    else /* Edge keys[index] doesn't exist */
    {
      return null pointer
    }
  }
  else { return root } /* Terminal condition */
}

```

or $f \in \text{St}(s)$ must also be removed along with any boundary/coboundary relations of f . The implemented removal algorithm traverses up the complex and removes simplex s and all cofaces of s level by level.

4.1.2 Searching and Traversing Topological Relations. The algorithms for retrieving a simplex as well as for basic traversals from one simplex to another across the data structure are the same. Given a starting simplex, and an array of keys up, the new simplex can be found recursively by Algorithm 2; the annotated code used is shown in Section A.3. Traversals from one simplex to another require a key lookup followed by a pointer dereference and therefore occur in approximately constant time ($O(1)$). Since all topological relations are stored, the traversal order across the array of keys does not matter. The same algorithm can be applied going down in dimension. For the retrieval of an arbitrary simplex, we start the search up from the root node of the complex.

4.2 Traversals

Thus far, we have presented algorithms for the creation of a simplicial complex as well as the basic traversal across faces and cofaces. For many applications, other traversals, such as by adjacency, may be more useful. We present several built-in traversal algorithms as well as the visitor design pattern for complicated operations.

4.2.1 By Level. It is often useful to have a traversal over all simplices of the same level. For example, iterating across all vertices to compute a center of mass. To support this in an efficient fashion, simplices of the same dimension are stored in a level specific map of iIDs to node pointers. Notably, the map for each level is instantiated with the correct user specified node type with respect to level at compile time. To achieve this, we again use variadic templates to generate a tuple of maps, where each tuple element corresponds to the map for a specific level's node type.

Since *asc_nodes* are templated on the integral level, we can use a template type map to map an integral sequence to the node pointer type,

$$\text{tuple}\langle 1, 2, 3, \dots \rangle \xrightarrow{\text{Node}\langle k \rangle *} \text{tuple}\langle \text{Node}\langle 1 \rangle *, \text{Node}\langle 2 \rangle *, \text{Node}\langle 3 \rangle *, \dots \rangle,$$

producing a tuple of integrally typed node pointers. Subsequently, we can map again to generate a tuple of maps,

$$\text{tuple}\langle \text{Node}\langle 1 \rangle *, \text{Node}\langle 2 \rangle *, \dots \rangle \xrightarrow{\text{map}\langle \text{int}, T \rangle} \text{tuple}\langle \text{map}\langle \text{int}, \text{Node}\langle 1 \rangle * \rangle, \text{map}\langle \text{int}, \text{Node}\langle 2 \rangle * \rangle, \dots \rangle.$$

By using this variadic template mapping strategy, we now have the correct typenames assigned. Any level of the tuple can be accessed by getting the integral level using functions in the C++ standard library. Variations of this mapping strategy are also used to construct the `SimplexSet` and `SimplexMap` structures below.

For end users, the implementation details are entirely abstracted away. Continuing from the example above, iteration over all vertices of simplicial complex, mesh, can be performed using the provided iterator adaptors as follows.

```
// Deduces the type of nid = ASC::SimplexID<1>
for(auto nid : mesh.get_level_id<1>()){
    std::cout << nid << std::endl;
}
```

Listing. 1. Example use of iterator adaptors for traversal across vertices of mesh.

The function `get_level_id<k>()` retrieves level k from the tuple and returns an iterable range across the corresponding map.

4.2.2 By Adjacency. Many geometric algorithms operate on the local neighborhood of a given simplex. Unlike other data structures such as the halfedge, CASC does not store the notion of the next simplex. Instead, adjacency is identified by searching for simplices with shared faces or cofaces in the complex. The algorithm for finding neighbors with shared faces is shown in Algorithm 3. We note that the set of simplices with shared faces may be different than the set of simplices with shared cofaces. Both adjacency definitions have been implemented, and we leave it to the end user to select the relevant function. Once a neighbor list has been aggregated, it can be traversed using standard methods. While the additional adjacency lookup step is extra in comparison to other data structures, in many cases, the generation of neighbor lists need only

ALGORITHM 3: Get the neighbors of a simplex, s , by inspecting the faces of s

Input: s : The simplex to get the neighbors of.

Output: List of neighbors

```
Function getNeighborsUp( $s$ ) /* Get the neighbors */
{
    Create an empty list of neighbors
    /* Follow all coboundary relations up from  $s$ . */
    for ( $a \in s.up$ )
    {
        SimplexID  $id = s \cup a$ 
        for ( $b \in id$ ) /* Go down from  $id$  */
        {
            if ( $(id \setminus b \neq s)$ ) { Add  $id \setminus b$  to neighbors } /* Do not add self to neighbors */
        }
    }
    return neighbors
}
```

be done once and cached. The tradeoff is that CASC offers facile manipulations of the topology without having to worry about reorganizing neighbor pointers.

4.2.3 Traversals Over Multiple Node Types. When performing more complicated traversals, such as iterating over the star of a simplex, multiple node types may be encountered. In order to avoid typename comparison-based branch statements, we have implemented visitor design pattern-based breadth first searches (BFS). The visitor design pattern refers to a double dispatch strategy where a traversal function takes a visitor functor, which implements an overloaded `visit()` function. At each node visited, the traversal function will call `visit()` on the current node. Since the functor overloads `visit()` per node type, the compiler can deduce which visit function to call. Example pseudocode is shown in Listing 2. This double dispatch strategy eliminates the need for extensive runtime typename comparisons and enables easy traversals over multiple node types. We provide breadth first traversals up and down the complex from a set of simplices. These visitor traversals are used extensively in the complex operations described below.

```
template <typename Complex>
struct Visitor{
    template <std::size_t k>
    using Simplex = typename Complex::template SimplexID<k>;
    // General template prototype
    template <std::size_t level>
    bool visit(Complex& F, Simplex<level> s){
        return true;
    }
    // Specialization for vertices
    bool visit(Complex& F, Simplex<1> s){
        s.position = s.position*2;
        return true;
    }
    // Specialization for faces
    bool visit(Complex& F, Simplex<3> s){
        s.color = 'green';
        return true;
    }
};

void BFS(ASC& mesh, Visitor&& v){
    // ... traversal code
    v.visit(mesh, currentSimplex);
    // NOTE: visit is overloaded and called based on function prototype.
}

void main(){
    // ... define simplicial complex traits for a surface mesh
    ASC mesh = ASC(); // construct the mesh object
    // ... insert some simplices etc.
    BFS(mesh, Visitor());
}
```

Listing. 2. Example pseudocode of double dispatch to traverse the complex while scaling the mesh by 2 and coloring the faces green.

4.3 Complex Operations

4.3.1 Star/Closure/Link. The star, link, and closure can be computed using the visitor breadth first traversals to collect simplices. These operations typically produce a set of simplices spanning multiple simplex dimensions, and thus simplex typenames, which cannot be stored in a traditional C++ set. We have implemented a multi-set data structure called the `SimplexSet`, which is effectively a tuple of typed sets corresponding to each level. The `SimplexSet` is constructed using the same mapping strategy as the tuple of maps used for the iteration across levels. For convenience, we provide functions for typical set operations such as insertion, removal, search, union, intersection, and difference. Using a combination of the star and closure functions with `SimplexSet` difference, we can get the link by Equation (3).

4.3.2 Metadata Aware Decimation. We have implemented a general decimation algorithm that operates by collapsing simplices of higher dimensions into a vertex. While edge collapses for 2-manifolds are well studied, a general dimensional collapse is useful for decimating higher-dimensional meshes used to solve PDEs such as those encountered in general relativity. Since simplices are being removed from the complex, user data may be lost. Our implementation is metadata aware and allows the user to specify what data to keep post decimation. This is achieved by using a recursive algorithm to produce a map of removed simplices to new simplices. The user can use this mapping to define a function that maps the original stored data to the post decimation topology. This decimation strategy is implemented as an inplace operation yielding decimated mesh containing data mapped according to a user specified callback function.

This decimation algorithm is a generalization of an edge collapse operation to arbitrary dimensions. It is formally defined as follows:

Definition 4.1. Given simplicial complex $\mathcal{F}$, simplex to decimate $s \in \mathcal{F}$, vertex set V of $\mathcal{F}$, and new vertex $p \notin V$, we define function,

$$\varphi(f) = \begin{cases} f & \text{if } f \cap s = \emptyset \\ p \cup (f \setminus s) & \text{if } f \cap s \neq \emptyset \end{cases}, \quad (6)$$

where f is any simplex $f \in \mathcal{F}$. We define the decimation of $\mathcal{F}$ by replacing s with p as $\varphi(\mathcal{F})$.

Note that decimation under this definition is not guaranteed to preserve the topology, as can be seen by decimating any edge or face of a tetrahedron.

Decimation of a simplicial complex must result in a valid triangulation. Here, we show that decimation by Definition 4.1 produces a valid abstract simplicial complex.

THEOREM 4.2. $\varphi(\mathcal{F})$ is an abstract simplicial complex.

PROOF. Given simplices y and x , let $y \in \varphi(\mathcal{F})$ and let $x \subset y$. We will show that $x \in \varphi(\mathcal{F})$. There are two cases for $y \cap p$, as they can be disjoint or intersecting.

Considering the disjoint case where $y \cap p = \emptyset$. This implies that $y \in \mathcal{F}$ and $y \cap s = \emptyset$. Since $\mathcal{F}$ is a simplicial complex, $y \in \mathcal{F}$ implies that $x \in \mathcal{F}$. Furthermore, since $y \cap s = \emptyset$ and $x \subset y$, then $x \cap s = \emptyset$ implying that $\varphi(x) = x$; thus, $x \in \varphi(\mathcal{F})$.

Alternately, in the intersecting case where $y \cap p \neq \emptyset$, that is $y \cap p = p$. Since $x \subset y$, there are two sub-cases where x either contains p or does not.

Supposing that $x \cap p = \emptyset$. Then, $x \subset (f \setminus s)$ implying that $x \in \mathcal{F}$ and $x \cap s = \emptyset$. Therefore, by the disjoint case, $\varphi(x) = x$ implies that $x \in \varphi(\mathcal{F})$.

Supposing that $x \cap p = p$. We can rewrite any such simplex x as $x = w \cup p$ where $w \cap p = \emptyset$. Furthermore, we can write that $y = x \cup r$ where $x \cap r = \emptyset$ and $r \cap p = \emptyset$; thus, $y = w \cup p \cup r$. There exists some set q such that $f = w \cup r \cup q$ and $q \subseteq s$ such that $\varphi(f) = p \cup ((w \cup r \cup q) \setminus s) = y$.

(a) A geometric realization of the complex before and after decimation.

(b) Explicitly drawn out Hasse diagrams for the constructed example where the TOP is before, and BOTTOM is after decimation. Grey arrows mark the relationships between sets of simplices before and after. Because there is always a mapping, users can define strategies to manage the stored data.

Fig. 5. The example decimation of edge $s = \{3, 4\}$ in a constructed example.

Since $f \in \mathcal{F}$, then $w \cup q \in \mathcal{F}$. Therefore, $\varphi(w \cup q) = p \cup ((w \cup q) \setminus s) = p \cup w = x$ and thus $x \in \varphi(\mathcal{F})$.

For all cases and sub-cases, we have shown that $x \in \varphi(\mathcal{F})$; therefore, $\varphi(\mathcal{F})$ is an abstract simplicial complex. $\square$

A pseudocode implementation for this decimation is provided in Algorithm 4. Given some simplicial complex $\mathcal{F}$ and simplex $s \in \mathcal{F}$ to decimate, this algorithm works in four steps. First, we compute the complete neighborhood, $nbhd = \text{St}(\text{Cl}(s))$, of s . Simplices not in the complete neighborhood will be invariant under φ and are ignored. Next, we use a nested set of breadth first searches to walk over the complete neighborhood and compute $p \cup f \setminus s$ for each simplex in the neighborhood. The results are inserted into a SimplexMap, which maps $\varphi(f)$ to a SimplexSet of all f that map to $\varphi(f)$. Third, we iterate over the SimplexMap and run the user defined callback on each mapping to generate a list of new simplices and associated mapped data stored in SimplexDataSet. Finally, the algorithm removes all simplices in the complete neighborhood and inserts the new mapped simplices.

An example application of this decimation operation is shown in Figure 5. In Figure 5(a), we show the geometric realization of the complex before and after the decimation of simplex $\{3, 4\}$. In Figure 5(b), we show the detailed Hasse diagrams for the constructed example. Note that there

ALGORITHM 4: Decimate a simplex by collapsing it into a vertex.**Input:** F : the simplicial complex; s : the simplex to decimate; $clbk$: a callback function to handle the data mapping**Output:** Simplicial complex with simplex collapsed according to Def. 4.1

```

Simplex np = F.newVertex() /* Create a dummy new vertex to map to */
SimplexSet N /* For the complete neighborhood */
SimplexDataSet data /* Data structure to store (simplex name, simplex data) pairs */
SimplexMap simplexMap /* Data structure to store New Simplex -> SimplexSet map */

/* Get the complete neighborhood (all simplices that are associated with s). */
for (vertex  $v$  of  $s$ )
{
    BFSup (simplex  $i \in St(v)$ )
    {
        if ( $j \notin N$ ) { N.insert ( $j$ ) }
    }
}

/* Backup the complete neighborhood. These simplices will be destroyed eventually. */
SimplexSet doomed = N
/* Generate the before-after mapping */
BFSdown (simplex  $i \in Cl(s)$ ) /* MainVisitor */
{
    BFSup (simplex  $j \in N \cap St(i)$ ) /* InnerVisitor */
    {
        /*  $i$  maps to  $np$  so we need to connect  $j$  to  $np$  instead */
        Name newName =  $np \cup j \setminus i$ 
        SimplexSet grab
        BFSdown (simplex  $k \in N \cap Cl(j)$ ) /* GrabVisitor */
        {
            /* Grab dependent simplices that have not been grabbed yet. Grabbed simplices
            will map to simplex  $newName$  */
            N.remove( $k$ )
            grab.insert( $k$ )
        }
        if ( $newName \notin simplexMap$ ) { simplexMap.insert(pair( $newName$ , grab)) }
        else { simplexMap[ $newName$ ].insert( $grab$ ) }
    }
}

for ({ $newName$ ,  $grabbed$ }  $\in simplexMap$ )
{
    /* Run the user's callback to map simplex data. */
    DataType mappedData = (*clbk)(F, name( $j$ ), newName, grabbed)
    data.insert({ $newName$ , mappedData}) /* Insert a pair containing new simplex name and
    data */
}

/* Iterate over the complete neighborhood and remove simplices */
performRemoval( $F$ , doomed)
/* Iterate over data and append mapped simplices and data */
performInsertion( $F$ , data)

```

Table 1. Traversal Order of the Visitors for the Decimation
Shown in Figure 5(a)

Order	MainVisitor	InnerVisitor	GrabVisitor	Maps to
1	{3,4}	{3,4}	{3,4}, {3}, {4}	{6}
2	{3,4}	{1,3,4}	{1,3,4}, {1,3}, {1,4}	{1,6}
3	{3,4}	{3,4,5}	{3,4,5}, {3,5}, {4,5}	{3,6}
4	{3}	{0,3}	{0,3}	{0,6}
5	{3}	{0,1,3}	{0,1,3}	{0,1,6}
6	{3}	{0,3,5}	{0,3,5}	{0,5,6}
7	{4}	{2,4}	{3,5}	{5,6}
8	{4}	{1,2,4}	{1,2,4}	{1,2,6}
9	{4}	{2,4,5}	{2,4,5}	{2,5,6}

are two possible mapping situations. In one case, $f \in \text{St}(\text{Cl}(s)) \cap \text{Cl}(\text{St}(s))$, groups of simplices are merged. In the other case, simplices $f \in \text{Cl}(\text{St}(s)) \setminus \text{St}(\text{Cl}(s))$ only need to be reconnected to the new merged simplices. By carefully choosing the traversal order, some optimizations can be made.

We apply the decimation on the constructed example shown in Figure 5 and show the order of operations with respect to the current visited simplex for each visitor function in Table 1. Starting out, MainVisitor and InnerVisitor will visit {3,4}. At this point, GrabVisitor will search BFS down from {3,4} to grab the set {{3,4}, {3}, {4}} and remove it from the neighborhood, eliminating some future calculations at {3} and {4}. All simplices in this set will map to new simplex {6} after decimation. Continuing upward, InnerVisitor will find simplex {1,3,4} and GrabVisitor will grab set {{1,3,4}, {1,3}, {1,4}}. Again, this set of simplices will map to common simplex {1,6} post decimation. A similar case occurs with simplex {3,4,5}. At this point, all simplices in $\text{St}(\text{Cl}(s)) \cap \text{Cl}(\text{St}(s))$ have been visited and removed from the neighborhood and MainVisitor continues BFS down, finds {3}, and calls BFS up (InnerVisitor). Note that since simplex {3} has already been grabbed, InnerVisitor will continue upward and find {0,3}. Looking down, there are no simplices that are faces of {0,3} in the neighborhood. So on and so forth.

To reiterate, GrabVisitor grabs the set of simplices that will be mapped to a common simplex. We show here that the order in which simplices are grabbed by Algorithm 4 will preserve that all simplices $f = w \cup q$ where $q \subseteq s$ will map to $\phi(f) = w \cup p$.

When visiting any simplex $f \supseteq q$ where $q \subseteq s$ and q corresponds to simplices visited by MainVisitor, we can write f as $f = w \cup q$ where the sets w and q are disjoint. Looking down from f , all simplices fall into two cases: $g = v \cup q$ where $\emptyset \subseteq v \subsetneq w$, or $h = w \cup t$ where $t \subsetneq q$. All simplices of form g , at worst case, will be grabbed while InnerVisitor proceeded BFS up from q . Remaining simplices h can be grouped with f and correctly mapped to $w \cup p$.

We note that in some non-manifold cases, GrabVisitor will not always grab set members in one visit. Supposing that we removed simplex {1,3,4} from the constructed example, in this case, InnerVisitor cannot visit {1,3,4} and simplices {1,3} and {1,4} will not be grouped. Instead, {1,3} and {1,4} will be found individually when MainVisitor visits {3} then {4}. To catch this case and correctly map {1,3} and {1,4} to {1,6}, we use a SimplexMap to aggregate all maps prior to proceeding. We note that in all cases starting with a valid simplicial complex, this implementation of the general collapse of simplex s visits each member in $\text{St}(\text{Cl}(s))$ and maps according to Def. 4.1, producing a valid simplicial complex. There is no guarantee that the result will have the same topological type as the pre-decimated mesh. The preservation of the topological type under decimation is often a desirable trait. We will show how to verify the Link Condition for when edge collapse of a 2-manifold will preserve the topological type in Section 5.1.

5 SURFACE MESH APPLICATION EXAMPLE

CASC is a general simplicial complex data structure that is suitable for use in mesh manipulation and processing. For example, we can use CASC as the underlying representation for an orientable surface mesh. Using a predefined `Vertex` class that is wrapped around a tensor library, and a class `Orientable` that wraps an integer, we can easily create a surface mesh embedded in $\mathbb{R}^3$.

```
using Vector = tensor<double,3,1>;
struct Vertex {
    Vector position;
    // ... other helpful vertex functions;
};
struct Orientable {
    int orientation;
};
struct complex_traits
{
    using KeyType = int;
    using NodeTypes = util::type_holder<void,Vertex,void,Orientable>;
    using EdgeTypes = util::type_holder<Orientable,Orientable,Orientable>;
};
using SurfaceMesh = simplicial_complex<complex_traits>;
```

In this case, 1-simplices will store a `Vertex` type while faces and all edges will store `Orientable`. Using `SurfaceMesh`, we can easily create functions to load or write common mesh filetypes such as OFF as shown in the included library examples.

ALGORITHM 5: Define the orientation of a topological relation.

Input: Simplices a and b where $b = a \cup v$ and v is a vertex.

Output: The orientation of edge $a \rightarrow b$

```
int orient = 1
for (Vertex u : a) /* For each vertex u in simplex a */
{
    if (v > u) { orient *= -1; }
}
return orient;
```

We can define a boundary morphism that applies on an ordered k -simplex,

$$\partial_i^k([a_0, \dots, a_{k-1}]) = (-1)^i([a_0, \dots, a_{k-1}] \setminus \{a_i\}), \quad (7)$$

where $a_i < a_{i+1}$. Using Algorithm 5, we can apply this morphism to assign a ± 1 orientation to each topological relation in the complex. Subsequently, for orientable manifolds, we can compute orientations of faces f_1 and f_2 that share edge e such that,

$$\text{Orient}(e_1) \cdot \text{Orient}(f_1) + \text{Orient}(e_2) \cdot \text{Orient}(f_2) = 0, \quad (8)$$

where e_1 and e_2 correspond to the edge up from e to f_1 and f_2 , respectively. Doing so, we create an oriented simplicial complex.

Supposing that we wish to compute the tangent of a vertex as defined by the weighted average tangent of incident faces. This is equivalent to computing the oriented wedge products of each

incident face. This can be written generally as,

$$\text{Tangent}(v) = \frac{1}{N} \sum_{i=0}^N \text{Orient}(f_i) \cdot (\partial_j(f_i) \wedge \partial_k(f_i)) \quad (9)$$

$$= \frac{1}{N} \sum_{i=0}^N \text{Orient}(f_i) \cdot \frac{1}{2} (e_{i,j} \otimes e_{i,k} - e_{i,k} \otimes e_{i,j}), \quad (10)$$

where N is the number of incident faces, f_i is incident face i , j and k are indices of vertex members of f_i not equal to v , and $e_{i,j} = \partial_j(f_i)$. This can be easily computed using a templated recursive function.

```
// Terminal case
auto getTangentH(const SurfaceMesh& mesh,
    const Vector& origin,
    SurfaceMesh::SimplexID<SurfaceMesh::topLevel> curr){
    return (*curr).orientation;
}

template <std::size_t level, std::size_t dimension>
auto getTangentH(const SurfaceMesh& mesh,
    const Vector& origin,
    SurfaceMesh::SimplexID<level> curr){
    tensor<double, 3, SurfaceMesh::topLevel - level> rval;
    auto cover = mesh.get_cover(curr); // Lookup coboundary relations
    for(auto v : cover){
        auto edge = *mesh.get_edge_up(curr, v); // Get the edge object
        const auto& vtx = (*mesh.get_simplex_up({v})).position; // Vertex v
        auto next = mesh.get_simplex_up(curr, v); // Simplex curr union v
        rval += edge.orientation * (v-origin) * getTangentH(mesh, origin, next);
    }
    return rval/cover.size();
}
```

This demonstrates the ease using the CASC library as an underlying simplicial complex representation. Using the provided Application Programming Interface (API), it is easy to traverse the complex to perform any computations.

5.1 Preservation of Topology Type of Surface Mesh under Edge Decimation by Contraction

Suppose we wish to decimate a surface mesh by edge contraction under Def. 4.1 without changing the topology of the complex. This can be verified by checking the Link Condition, defined with proof from Edelsbrunner [12], stating,

LEMMA 5.1. *Let F be a triangulation of a 2-manifold. The contraction of $ab \in F$ preserves the topological type if and only if $\text{Lk}(a) \cap \text{Lk}(b) = \text{Lk}(ab)$.*

Revisiting the example from Figure 5(a), we can construct the topology and check the Link Condition using operations supported by the CASC library.

```
// Construct the topology
SurfaceMesh mesh;
mesh.insert({0,1,3});
mesh.insert({0,3,5});
mesh.insert({1,3,4});
```

```

mesh.insert({3,4,5});
mesh.insert({1,2,4});
mesh.insert({2,4,5});

// Get simplices and edge
SimplexSet<SurfaceMesh> A,B,AB, AcapB;
auto ab = mesh.get_simplex_up({3,4});
auto a = mesh.get_simplex_up({3});
auto b = mesh.get_simplex_up({4});

// Compute the links of each
getLink(mesh, a, A);
getLink(mesh, b, B);
getLink(mesh, ab, AB);

// Link(a): Simplices {1}, {4}, {5}, {1,4}, {0,1}, {0,5}, {4,5}
std::cout << A << std::endl;
// Link(b): Simplices {5}, {2}, {3}, {1}, {3,5}, {1,3}, {1,2}, {2,5}
std::cout << B << std::endl;
// Link(AB): Simplices {1}, {5}
std::cout << AB << std::endl;

set_intersection(A, B, AcapB);
// Link(a) \cap Link(b): Simplices {1}, {5}
std::cout << AcapB << std::endl;

// Check the Link Condition
std::cout << (AB == AcapB) << std::endl; // Evaluates to true

```

The `getLink()` function utilizes a series of visitor breadth first traversals down then up the complex, collecting the set of simplex into a `SimplexSet` object. Then, using set operations provided by `SimplexSet`, the set difference and equality comparison are performed. This example highlights the simplicity, clarity, and transparency of using the CASC library.

6 CONCLUSIONS

CASC provides a general simplicial complex data structure, which allows the storage of user defined types at each simplex level. The library comes with a full-featured API providing common simplicial complex operations, as well as support for complex traversals using a visitor. We also provide a metadata-aware decimation algorithm, which allows users to collapse simplices of any dimension while preserving data according to a user-defined mapping function. Our implementation of CASC using a strongly-typed language is only possible due to recent innovations in language tools. The CASC API abstracts away most of the complicated templating, allowing it to be both modern and easy to use. We anticipate that CASC will not only be of use for the ET community, but microscopy and modeling as a whole along with other fields like applied mathematics and Computer-Aided Design (CAD).

One limitation is the ease of extending to other languages. The generality of CASC is reliant upon the C++ compiler. Sacrificing this, specific realizations of CASC can be wrapped using tools like Simplified Wrapper and Interface Generator (SWIG) for use in other languages. Another limitation of CASC is the memory efficiency. The current Hasse diagram-based implementation was selected for the sake of transparency, ease of traversal, and manipulation. Optimizations to the memory efficiency of CASC can be made by employing more compact representations. The vari-

adic template approach we use to attach user data to simplices is compatible with data structures, which explicitly represent all simplices but only a subset of topological relations. This includes data structures such as SimplexTree [6], IS [10], and SIG [8] among others. Other compressed data structures that skip levels of low importance using implicit nodes are not compatible with the current CASC implementation. Skipped levels would need to be implemented as exceptions to the combinatorial variadic rules. Similarly, although CASC in its current form is restricted to the representation of simplicial complexes, the combinatorial strategy can be easily adapted to support other regular polytopes by changing the boundary relation storage rules. In the future, we hope to incorporate parallelism into the CASC library. A copy of CASC along with online documentation can found on GitHub: <https://github.com/ctlee/casc>.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees for their valuable comments and helpful suggestions. R. E. Amaro is a cofounder and scientific advisor of, and has equity interest in, Actavalon, Inc.

REFERENCES

- [1] [n.d.]. CGAL, Computational Geometry Algorithms Library. Retrieved from <http://www.cgal.org>.
- [2] I. Babuška and A. K. Aziz. 1976. On the angle condition in the finite element method. *SIAM J. Numer. Anal.* 13, 2 (Apr. 1976), 214–226. DOI: <https://doi.org/10.1137/0713021>
- [3] Randolph E. Bank, Andrew H. Sherman, and Alan Weiser. 1983. Refinement algorithms and data structures for regular local mesh refinement. In *Sci. Comput. Appl. Mathematics*. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.40.2544>.
- [4] Randolph E. Bank and Jinchao Xu. 1996. An algorithm for coarsening unstructured meshes. *Numer. Math.* 73, 1 (Mar. 1996), 1–36. DOI: <https://doi.org/10.1007/s002110050181>
- [5] Jean-Daniel Boissonnat, Olivier Devillers, and Samuel Hornus. 2009. Incremental construction of the delaunay graph in medium dimension. In *Annu. Symp. Comput. Geom.* Aarhus, Denmark, 208–216. <https://hal.inria.fr/inria-00412437>.
- [6] Jean-Daniel Boissonnat and Clément Maria. 2014. The simplex tree: An efficient data structure for general simplicial complexes. *Algorithmica* 70, 3 (Nov. 2014), 406–427. DOI: <https://doi.org/10.1007/s00453-014-9887-3>
- [7] David Canino, Leila De Floriani, and Kenneth Weiss. 2011. IA*: An adjacency-based representation for non-manifold simplicial shapes in arbitrary dimensions. *Comput. Graph.* 35, 3 (June 2011), 747–753. DOI: <https://doi.org/10.1016/j.cag.2011.03.009>
- [8] Leila De Floriani, David Greenfieldboyce, and Annie Hui. 2004. A data structure for non-manifold simplicial D - complexes. In *Proc. 2004 Eurographics/ACM SIGGRAPH Symp. Geom. Process. (SGP'04)*. ACM Press, New York, New York, 83. DOI: <https://doi.org/10.1145/1057432.1057444>
- [9] Leila De Floriani and Annie Hui. 2005. Data structures for simplicial complexes: An analysis and a comparison. In *Proc. 3rd Eurographics Symp. Geom. Process.* Eurographics Association, Vienna, 119. <http://dl.acm.org/citation.cfm?id=1281920.1281940>.
- [10] Leila De Floriani, Annie Hui, Daniele Panozzo, and David Canino. 2010. A dimension-independent data structure for simplicial complexes. In *Proc. 19th Int. Meshing Roundtable*. Springer Berlin, Germany, 403–420. DOI: https://doi.org/10.1007/978-3-642-15414-0_24
- [11] Vladimir Dyedov, Navamita Ray, Daniel Einstein, Xiangmin Jiao, and Timothy J. Tautges. 2015. AHF: Array-based half-facet data structure for mixed-dimensional and non-manifold meshes. *Eng. Comput.* 31, 3 (July 2015), 389–404. DOI: <https://doi.org/10.1007/s00366-014-0378-6>
- [12] Herbert Edelsbrunner. 2001. *Geometry and Topology for Mesh Generation*. Cambridge University Press, Cambridge. DOI: <https://doi.org/10.1017/CBO9780511530067>
- [13] Zhanheng Gao, Zeyun Yu, and Michael Holst. 2013. Feature-preserving surface mesh smoothing via suboptimal Delaunay triangulation. *Graph. Models* 75, 1 (Jan. 2013), 23–38. DOI: <https://doi.org/10.1016/j.gmod.2012.10.007>
- [14] M. Holst, O. Sarbach, M. Tiglio, and M. Vallisneri. 2016. The emergence of gravitational wave science: 100 years of development of mathematical theory, detectors, numerical algorithms, and data analysis tools. *Bull. Amer. Math. Soc.* 53 (2016), 513–554. <http://dx.doi.org/10.1090/bull/1544>.
- [15] M. Holst and C. Tiee. 2018. Finite element exterior calculus for parabolic evolution problems on Riemannian hypersurfaces. *J. Comput. Math.* 36, 6 (2018), 792–832.

- [16] A. Liu and B. Joe. 1994. Relationship between tetrahedron shape measures. *BIT* 34, 2 (June 1994), 268–287. DOI : <https://doi.org/10.1007/BF01955874>
- [17] Navamita Ray, Iulian Grindeanu, Xinglin Zhao, Vijay Mahadevan, and Xiangmin Jiao. 2015. Array-based hierarchical mesh generation in parallel. *Procedia Eng.* 124 (2015), 291–303. DOI : <https://doi.org/10.1016/j.proeng.2015.10.140>
- [18] T. Regge. 1961. General relativity without coordinates. *Nuovo Cim.* 19, 3 (Feb. 1961), 558–571. DOI : <https://doi.org/10.1007/BF02733251>
- [19] Elijah Roberts. 2014. Cellular and molecular structure as a unifying framework for whole-cell modeling. *Curr. Opin. Struct. Biol.* 25 (Apr. 2014), 86–91. DOI : <https://doi.org/10.1016/j.sbi.2014.01.005>
- [20] Hang Si. 2015. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Trans. Math. Softw.* 41, 2 (Feb. 2015), 1–36. DOI : <https://doi.org/10.1145/2629697>
- [21] Eric Vanden-Eijnden and Maddalena Venturoli. 2009. Markovian milestoning with Voronoi tessellations. *J. Chem. Phys.* 130, 19 (May 2009), 194101. DOI : <https://doi.org/10.1063/1.3129843>
- [22] Zeyun Yu, Michael J. Holst, and J. Andrew McCammon. 2008. High-fidelity geometric modeling for biomedical applications. *Finite Elem. Anal. Des.* 44, 11 (July 2008), 715–723. DOI : <https://doi.org/10.1016/j.finel.2008.03.004>
- [23] Zeyun Yu, Michael J. Holst, Takeharu Hayashi, Chandrajit L. Bajaj, Mark H. Ellisman, J. Andrew McCammon, and Masahiko Hoshijima. 2008. Three-dimensional geometric modeling of membrane-bound organelles in ventricular myocytes: Bridging the gap between microscopic imaging and mathematical simulation. *J. Struct. Biol.* 164, 3 (Dec. 2008), 304–313. DOI : <https://doi.org/10.1016/j.jsb.2008.09.004>
- [24] Yongjie Zhang. 2016. *Geometric Modeling and Mesh Generation from Scanned Images*. Chapman and Hall/CRC. DOI : <https://doi.org/10.1201/b19466>

Received September 2017; revised November 2018; accepted March 2019