

# Efficient Representation of Computational Meshes

**A. Logg**

Simula Research Laboratory

Dept. of Informatics, University of Oslo, Norway

E-mail: logg@simula.no

**Abstract:** We present a simple yet general and efficient approach to representation of computational meshes. Meshes are represented as sets of *mesh entities* of different topological dimensions and their *incidence relations*. We discuss a straightforward and efficient storage scheme for such mesh representations and efficient algorithms for computation of arbitrary incidence relations from a given initial and minimal set of incidence relations. The general representation may harbor a wide range of computational meshes, and may also be specialized to provide simple user interfaces for particular meshes, including simplicial meshes in one, two and three space dimensions where the mesh entities correspond to vertices, edges, faces and cells. It is elaborated on how the proposed concepts and data structures may be used for assembly of variational forms in parallel over distributed finite element meshes. Benchmarks are presented to demonstrate efficiency in terms of CPU time and memory usage.

**Keywords:** Mesh, mesh representation, mesh algorithms, mesh entity, parallel assembly

**Reference** to this paper should be made as follows: A. Logg (2008) ‘Efficient Representation of Computational Meshes’. An early version of this manuscript was presented at MekIT07: Fourth national conference on Computational Mechanics.

---

## 1 Introduction

---

The computational mesh is a central component of any software framework for the (mesh-based) solution of partial differential equations. To reduce run-time and enable the solution of large problems, it is therefore important that the computational mesh may be represented efficiently, both in terms of the speed of operations on the mesh or access of mesh data, and in terms of the memory usage for storing any given mesh in memory.

It is furthermore important that the data structure for the representation of the mesh is general enough to harbor a wide range of computational meshes. This generality must also be reflected in the programming interface to the mesh representation, to allow the implementation of general algorithms on the computational mesh. Many algorithms, such as the assembly of a linear system from a finite element variational problem may be implemented similarly for simplicial, quadrilateral and hexahedral meshes if the programming interface to the mesh representation does not enforce a specific interface limited to a specific mesh type. For example, if the entities on the boundary of a mesh (the *facets*) may be accessed in a similar way independently of the mesh dimension and not as *edges* in two space dimensions and *faces* in three space dimensions, one may use the

same code to apply boundary conditions in 2D and 3D.

In Karpeev and Knepley (2005), a very general and flexible representation of computational meshes is presented. The mesh is represented as a *sieve*, which is in general a directed acyclic graph with the mesh entities as points and directed edges describing how the mesh entities are connected. In this paper, we take a slightly less general approach but build on some of the concepts from Karpeev and Knepley (2005). In particular, we will represent the mesh as a set of *mesh entities* (corresponding to the *points* of the sieve) and their *incidence relations*. We also acknowledge the works Berti (2002, 2006), where similar concepts are defined and where the importance of *mesh iterators* for expressing generic algorithms on computational meshes is advocated.

The data structures and algorithms discussed in this paper have been implemented as a C++ library and is distributed as part of DOLFIN, see Hoffman et al. (2006). DOLFIN is a problem-solving environment for ordinary and partial differential equations and is developed as part of the FEniCS project for the automation of computational mathematical modeling, see FEniCS (2008); Dupont et al. (2003). Interfaces to DOLFIN are available in the form of

Copyright © 200x Inderscience Enterprises Ltd.

a C++ and a Python class library.

## 1.1 Design Goals

When designing the mesh library, we had the following design goals in mind for the mesh representation and its interface. The mesh representation should be *simple*, meaning that the data is represented in terms of basic C++ arrays `unsigned int*` and `double*`; it should be *generic*, meaning that it should not be specialized to say simplicial meshes in one, two and three space dimensions; and it should be *efficient*, meaning that operations on the mesh or access of mesh data should be fast and the storage should require minimal memory usage for any given mesh. Furthermore, the programming interface to the mesh representation should be *intuitive*, meaning that suitable abstractions (classes) should be available, including specialized interfaces for specific types of meshes as well as generic interfaces that enable dimension-independent programming; and it should be *efficient*, meaning that the overhead of the object-oriented interface should be minimized.

## 1.2 Outline

In the following section, we present the basic concepts that define the mesh representation and its interface. We then discuss the data structures of the C++ implementation of the mesh representation in DOLFIN, followed by a discussion of the algorithms used in DOLFIN to compute any given incidence relation from a given minimal set of incidence relations. Next, we demonstrate the programming interface to the mesh library. This is followed by a discussion of distributed (parallel) mesh data structures. Finally, we present a series of benchmarks to demonstrate the efficiency of the mesh representation and its implementation followed by some concluding remarks.

---

## 2 Concepts

---

The mesh representation is based on the following basic concepts: *mesh*, *mesh topology*, *mesh geometry*, *mesh entity* and *mesh connectivity*. Each of these concepts is mapped directly to the corresponding component (class) of the implementation.

A mesh is defined by its topology and its geometry. The mesh topology defines how the mesh is composed of its parts (the mesh entities) and the mesh geometry describes how the mesh is embedded in some metric space, typically  $\mathbb{R}^n$  for  $n = 1, 2, 3$ . A mesh topology (Figure 1) may be specified as a set of mesh entities (the vertices, edges etc.) and their connectivity (incidence relations). Different embeddings (geometries) may be imposed on any given mesh topology to create different meshes, e.g., when moving the vertices of a mesh in an ALE computation. Below, we discuss the two basic concepts mesh entity and mesh connectivity in some detail and also introduce the concept *mesh function*.




Figure 1: A mesh topology is a set of mesh entities (vertices, edges, etc.) and their connectivity (incidence relations), that is, which entities are connected (incident) to which entities.

### 2.1 Mesh Entities

A mesh entity is a pair  $(d, i)$ , where  $d$  is the topological dimension of the mesh entity and where  $i$  is a unique index for the mesh entity within its topological dimension, ranging from 0 to  $N_d - 1$  with  $N_d$  the number of entities of topological dimension  $d$ . We let  $D$  denote the maximal topological dimension over the mesh entities and set the topological dimension of the mesh equal to  $D$ . This is illustrated in Figure 2, where each mesh entity is labeled by its topological dimension and index  $(d, i)$ .




Figure 2: Each mesh entity of a mesh is identified with a pair  $(d, i)$ , where  $d$  is the topological dimension of the mesh entity and where  $i$  is a unique index for the mesh entity within its topological dimension, ranging from 0 to  $N_d - 1$  with  $N_d$  the number of entities of topological dimension  $d$ .

For convenience, we also name common entities of low topological dimension or codimension. We refer to entities of topological dimension 0 as *vertices*, entities of dimension 1 as *edges*, entities of dimension 2 as *faces*, entities of codimension 1 (dimension  $D - 1$ ) as *facets* and entities of codimension 0 (dimension  $D$ ) as *cells*. Thus, for a triangular mesh, the edges are also facets and the faces are also

cells, and for a tetrahedral mesh, the faces are also facets. This is summarized in Table 1.

Entity	Dimension	Codimension
Vertex	0	$D$
Edge	1	$D - 1$
Face	2	$D - 2$
Facet	$D - 1$	1
Cell	$D$	0

Table 1: Named entities of low topological dimension or codimension.

## 2.2 Mesh Connectivity

We refer to the set of incidence relations on a set of mesh entities as the *connectivity* of the mesh. For a mesh of topological dimension  $D$ , there are  $(D+1)^2$  different classes of incidence relations (connectivities) to consider. Each such class is denoted here by  $d \rightarrow d'$  for  $0 \leq d, d' \leq D$ . For any given mesh entity  $(d, i)$ , its connectivity  $(d \rightarrow d')_i$  is given by the set of incident mesh entities of dimension  $d'$ .

Thus, for a triangular mesh (of topological dimension  $D = 2$ ), there are nine different incidence relations of interest between the entities of the mesh. These are in turn  $0 \rightarrow 0$  (the vertices incident to each vertex),  $0 \rightarrow 1$  (the edges incident to each vertex),  $\dots$ ,  $D \rightarrow D$  (the cells incident to each cell).

For  $d > d'$ , the definition of incidence is evident. Mesh entity  $(d', i')$  is incident to mesh entity  $(d, i)$  if  $(d', i')$  is *contained* in  $(d, i)$ . Thus, the three vertices of a triangular cell form the set of incident vertices and the three edges form the set of incident edges. For  $d < d'$ , we define mesh entity  $(d', i')$  as incident to mesh entity  $(d, i)$  if  $(d, i)$  is incident to  $(d', i')$ . It thus remains to define incidence for  $d = d'$ . For  $d, d' > 0$ , we say that mesh entity  $(d', i')$  is incident to mesh entity  $(d, i)$  if both are incident to a common vertex, that is, a mesh entity of dimension zero, while for  $d = d' = 0$ , we say that  $(d', i')$  is incident to  $(d, i)$  if both are incident to a common cell, that is, a mesh entity of dimension  $D$ .

Together, the set of mesh entities and the connectivity (incidence relations) define the topology of the mesh. Note that the complete set of incidence relations  $d \rightarrow d'$  for  $0 \leq d, d' \leq D$  may be determined from the single class of incidence relations  $D \rightarrow 0$ , that is, the vertices of each cell in the mesh. We return to this below when we present an algorithm for computing any given class of incidence relations from the minimal set of incidence relations  $D \rightarrow 0$ .

## 2.3 Mesh Functions

We define a *mesh function* as a discrete function that takes a value on the set of mesh entities of a given fixed dimension  $0 \leq d \leq D$ . Mesh functions are simple objects but very useful. A real-valued mesh function may for

example be used to describe material parameters on the cells of a mesh. A boolean-valued mesh function may be used to set markers on cells or edges for adaptive refinement. Integer-valued mesh functions may be used to express inter-connectivity between two separate meshes. A typical use is when a boundary mesh is extracted from a given mesh (by identifying the set of facets that are incident to exactly one cell). One may then use a mesh function to describe the mapping from the cells in the extracted boundary mesh (which has topological dimension  $D - 1$ ) to the corresponding facets in the original mesh (which has topological dimension  $D$ ). Note that mesh functions are discrete and are not meant to represent for example a piecewise polynomial finite element function on the mesh.

---

## 3 Data Structures

---

The mesh representation as described in the previous section has been implemented as a small C++ class library and is available freely as part of the DOLFIN C++ finite element library, version 0.6.3 or higher. Each of the basic concepts *mesh*, *mesh topology*, *mesh geometry*, *mesh entity*, *mesh connectivity* and *mesh function* is realized by the corresponding class `Mesh`, `MeshTopology`, `MeshGeometry`, `MeshEntity`, `MeshConnectivity` and `MeshFunction`. All basic data structures are stored as static arrays of unsigned integers (`unsigned int*`) or floating point values (`double*`), which minimizes the cost of storing the mesh data and allows for quick access of mesh data. We discuss each of these classes/data structures in detail below.

### 3.1 The Class Mesh

The class `Mesh` stores a `MeshTopology` and a `MeshGeometry` that together define the mesh. The `MeshTopology` and `MeshGeometry` are independent of each other and of the `Mesh`. Although it is possible to work with the `MeshTopology` and `MeshGeometry` separately, they are most conveniently accessed through a `Mesh` class that holds a pair of a matching topology and geometry.

### 3.2 The Class MeshTopology

The class `MeshTopology` stores the topology of a mesh as a set of mesh entities and connectivities. For each pair of topological dimensions  $(d, d')$ ,  $0 \leq d, d' \leq D$ , the class `MeshTopology` stores a `MeshConnectivity` representing the set of incidence relations  $d \rightarrow d'$ . The mesh entities themselves need not be stored explicitly; they are stored implicitly for each topological dimension  $d$  as the set of pairs  $(d, i)$  for  $0 \leq i < N_d$ , where  $N_d$  is the number of mesh entities of topological dimension  $d$ . Thus, for each topological dimension, the class `MeshTopology` stores an (unsigned) integer  $N_d$ , from which the set of mesh entities  $\{(d, 0), (d, 1), \dots, (d, N_d - 1)\}$  may be generated.

### 3.3 The Class MeshGeometry

The class `MeshGeometry` stores the geometry of a mesh. Currently, only the simplest possible representation has been implemented, where only the coordinates of each vertex are stored. These coordinates are stored in a contiguous array `coordinates` of size  $nN_0$ , where  $n$  is the geometric dimension and  $N_0$  is the number of vertices.

### 3.4 The Class MeshEntity

The class `MeshEntity` provides a *view* of a given mesh entity  $(d, i)$ . The mesh entities themselves are not stored, but a `MeshEntity` may be generated from a given pair  $(d, i)$ . The class `MeshEntity` provides a convenient interface for accessing mesh data, in particular in combination with the concept of mesh iterators, as will be discussed in more detail below. Thus, one may for any given `MeshEntity` access its topological dimension  $d$ , its index  $i$  and its set of incidence relations (connected mesh entities) of any given topological dimension  $d'$ . Specialized interfaces are provided for the named mesh entities of Table 1 in the form of the following sub classes of `MeshEntity`: `Vertex`, `Edge`, `Face`, `Facet` and `Cell`.

### 3.5 The Class MeshConnectivity

The class `MeshConnectivity` stores the set of incidence relations  $d \rightarrow d'$  for a fixed pair of topological dimensions  $(d, d')$ . The set of incidence relations is stored as a contiguous `unsigned int` array `indices` of entity indices for dimension  $d'$  entities, together with an auxiliary `unsigned int` array `offsets` that specifies the offset into the first array for each entity of dimension  $d$ .<sup>1</sup> The size of the first array `indices` is equal to the total number of incident entities of dimension  $d'$  and the size of the second array `offsets` is equal to the total number of entities of dimension  $d$  plus one.

As an example, consider the storage of the set of incidence relations  $2 \rightarrow 0$ , that is the vertices of each cell, for the triangular mesh in Figure 3. The mesh has two entities of dimension  $d = 2$  and four entities of dimension  $d' = 0$ . Furthermore, each entity of dimension  $d = 2$  is incident to three entities of dimension  $d' = 0$ . The array `entities` is then given by `[0, 1, 3, 1, 2, 3]` and the array `offsets` is given by `[0, 3, 6]`.

### 3.6 The Class MeshFunction

The class `MeshFunction` stores a single array of  $N_d$  values on the mesh entities of a given fixed dimension  $d$ , and is templated over the value type. Typical uses include `MeshFunction<double>` for material parameters that take a constant value on each cell of a

<sup>1</sup>The storage is similar to the standard compressed row storage (CRS) format for sparse matrices, except that only the column indices need to be stored, not the values. Also note that the two arrays `indices` and `offsets` are private data structures of the class `MeshConnectivity`. The user is presented with a more intuitive interface, as will be demonstrated below.




Figure 3: The mesh connectivity  $2 \rightarrow 0$  (the vertices of each cell) for this triangular mesh with two cells and four vertices is stored as two arrays `indices = [0, 1, 3, 1, 2, 3]` and `offsets = [0, 3, 6]`.

mesh, `MeshFunction<bool>` for cell markers that indicate cells that should be refined, and `MeshFunction<unsigned int>` to store inter-mesh connectivity or sub domain markers.

### 3.7 Minimal Storage

The mesh data structures described above are summarized in Table 2. We note that the classes `Mesh` and `MeshTopology` function as “aggregate classes” that collect mesh data stored elsewhere, and that no data is stored in the class `MeshEntity`. All data is thus stored in the class `MeshConnectivity` (in the two arrays `indices` and `offsets`) and in the class `MeshGeometry` (in the array `coordinates`). Note that one `MeshConnectivity` object is stored for each pair of topological dimensions  $(d, d')$  for which the mesh connectivity has been initialized.

Data structure	Principal data
<code>Mesh</code>	<code>MeshTopology topology</code> <code>MeshGeometry geometry</code>
<code>MeshTopology</code>	<code>MeshConnectivity** connectivities</code>
<code>MeshGeometry</code>	<code>double* coordinates</code>
<code>MeshEntity</code>	–
<code>MeshConnectivity</code>	<code>unsigned int* indices</code> <code>unsigned int* offsets</code>

Principal data	Size
<code>MeshTopology topology</code>	–
<code>MeshGeometry geometry</code>	–
<code>MeshConnectivity** connectivities</code>	–
<code>double* coordinates</code>	$nN_0$
<code>unsigned int* indices</code>	$\mathcal{O}(N_d)$
<code>unsigned int* offsets</code>	$N_d + 1$

Table 2: Summary of mesh data structures.

As an illustration, consider the storage of a tetrahedral mesh with  $N_0$  vertices and  $N_3$  cells (tetrahedra) embedded in  $\mathbb{R}^3$  where we only store the set of incidence relations  $D \rightarrow 0$ . Each cell has four vertices, so the class `MeshConnectivity` stores  $4N_3 + N_3 + 1 \sim 5N_3$  values of type `unsigned int`. Furthermore, the class `MeshGeometry` stores  $3N_0$  values of type `double`. Thus, if an `unsigned int` is four bytes and a `double` is eight bytes, then the total size of the mesh is  $20N_3 + 24N_0$  bytes. For a standard uniform tetrahedral mesh of the unit square, the number of cells is approximately six times the number of vertices, so the total size of the mesh is

$$(20N_3 + 24N_0) \text{ b} = (20N_3 + 24N_3/6) \text{ b} = 24N_3 \text{ b}. \quad (1)$$

Thus, a mesh with 1,000,000 cells may be stored in just 24 Mb. Note that if additional mesh connectivity is computed, like the edges or facets of the tetrahedra, more memory will be required to store the mesh.

---

#### 4 Algorithms

---

In this section, we present the algorithms used by the DOLFIN mesh library to compute the mesh connectivity  $d \rightarrow d'$  for any given  $0 \leq d, d' \leq D$ . We assume that we are given an initial set of incidence relations  $D \rightarrow 0$ , that is, we know the vertices of each cell in the mesh.

The key to computing the mesh connectivities of a mesh is to compute the connectivities in a particular order. For example, if the vertices are known for each edge in the mesh ( $1 \rightarrow 0$ ), then it is straightforward to compute the edges incident to each vertex ( $0 \rightarrow 1$ ) as will be explained below. The computation is based on three algorithms that are used successively in a particular order to compute the desired connectivity. As a consequence, the computation of a certain connectivity  $d \rightarrow d'$  might require the computation of one or more other connectivities. We describe these algorithms in detail below. An overview is given in Figure 4

##### 4.1 Build

Algorithm 1 (Build) computes the connectivities  $D \rightarrow d$  and  $d \rightarrow 0$  from  $D \rightarrow 0$  and  $D \rightarrow D$  for  $0 < d < D$ . In other words, given the vertices and incident cells of each cell in the mesh, Algorithm 1 computes the entities of dimension  $d$  of each cell and for each such entity the vertices of that entity. Thus, if  $d = 1$ , then the edges of each cell and the vertices of each edge are computed.

The notation of Algorithm 1 requires some explanation. As before, we let  $(d \rightarrow d')_i$  denote the set of entities of dimension  $d'$  incident to entity  $(d, i)$ :

$$(d \rightarrow d')_i = \{(d', j) : (d', j) \text{ incident to } (d, i)\}. \quad (2)$$

Algorithm 1 also uses the operation

$$d \xrightarrow{\text{local}(D, i)} 0, \quad (3)$$




Figure 4: The three basic algorithms for computing connectivity. From the top: Build (computing connectivity  $D \rightarrow d$  and  $d \rightarrow 0$  from  $D \rightarrow 0$  and  $D \rightarrow D$ ), Transpose (computing connectivity  $d \rightarrow d'$  from  $d' \rightarrow d$ ) and Intersection (computing connectivity  $d \rightarrow d'$  from  $d \rightarrow d''$  and  $d'' \rightarrow d'$ ).



which denotes the set of vertex sets incident to the mesh entities of topological dimension  $d$  of a given cell  $(D, i)$ . To make this concrete, consider a triangular mesh (for which  $D = 2$ ) and take  $d = 1$ . If  $V_i = d \xrightarrow{\text{local}(D,i)} 0$ , then  $V_i$  denotes the set of vertex sets incident to the edges of triangle number  $i$ . The set  $V_i$  consists of three sets of vertices (one for each edge) and each set  $v_i \in V_i$  contains two vertices. In addition, Algorithm 1 uses the operation

$$\text{index}((D, j), d, v_i), \quad (4)$$

which denotes the index of the entity of dimension  $d$  in the cell  $(D, j)$  which is incident to the vertices  $v_i$ .

We may now summarize Algorithm 1 as follows. For each cell  $(D, i)$ , we create a set of candidate entities of dimension  $d$ , represented by their incident vertices in the set  $V_i$ . This operation is local on each cell and must be performed differently for each different type of mesh. We then iterate over each cell incident to the cell  $(D, i)$  and check for each candidate entity  $v_i \in V_i$  if it has already been created by any of the previously visited cells, making sure that two incident cells agree on the index of any common incident entity.

---

**Algorithm 1** Build( $d$ ), computing  $D \rightarrow d$  and  $d \rightarrow 0$  from  $D \rightarrow 0$  and  $D \rightarrow D$  for  $0 < d < D$

---

```

 $k = 0$ 
for each  $(D, i)$ 
   $V_i = d \xrightarrow{\text{local}(D,i)} 0$ 
  for each  $(D, j) \in (D \rightarrow D)_i$  such that  $j < i$ 
     $V_j = d \xrightarrow{\text{local}(D,j)} 0$ 
    for each  $v_i \in V_i$ 
      if  $v_i \in V_j$ 
         $l = \text{index}((D, j), d, v_i)$ 
         $(D \rightarrow d)_i = (D \rightarrow d)_i \cup (d, l)$ 
      else
         $(D \rightarrow d)_i = (D \rightarrow d)_i \cup (d, k)$ 
         $(d \rightarrow 0)_k = v_i$ 
         $k = k + 1$ 

```

---

## 4.2 Transpose

Algorithm 2 (Transpose) computes the connectivity  $d \rightarrow d'$  from the connectivity  $d' \rightarrow d$  for  $d < d'$ . For each entity of dimension  $d'$ , we iterate over the incident entities of dimension  $d$  and add the entities of dimension  $d'$  as incident entities to the entities of dimension  $d$ . We may thus compute for example the incident cells of each vertex (the cells to which the vertex belongs) by iterating over the cells of the mesh and for each cell over its incident vertices.

## 4.3 Intersection

Algorithm 3 (Intersection) computes the connectivity  $d \rightarrow d'$  from  $d \rightarrow d''$  and  $d'' \rightarrow d'$  for  $d \geq d'$ . For each entity  $(d, i)$  of dimension  $d$ , we iterate over each incident entity

---

**Algorithm 2** Transpose( $d, d'$ ), computing  $d \rightarrow d'$  from  $d' \rightarrow d$  for  $d < d'$

---

```

for each  $(d', j)$ 
  for each  $(d, i) \in (d' \rightarrow d)_j$ 
     $(d \rightarrow d')_i = (d \rightarrow d')_i \cup (d', j)$ 

```

---

$(d'', k)$  of dimension  $d''$  and for each such entity we iterate over each incident entity  $(d', j)$  of dimension  $d'$ . We then check if either  $(d, i)$  and  $(d', j)$  are entities of the same topological dimension or if  $(d', j)$  is completely contained in  $(d, i)$  by checking that each vertex incident to  $(d', j)$  is also incident to  $(d, i)$ , in which case  $(d', j)$  is added as an incident entity of entity  $(d, i)$ .

Here,  $d''$  must be chosen according to the definition of incidence given above. For example, we may take  $d'' = 0$  to compute the connectivity  $D \rightarrow D$  (the incident cells of each cell) by iterating over the vertices of each cell and for each such vertex iterate over the incident cells.

---

**Algorithm 3** Intersection( $d, d', d''$ ), computing  $d \rightarrow d'$  from  $d \rightarrow d''$  and  $d'' \rightarrow d'$  for  $d \geq d'$

---

```

for each  $(d, i)$ 
  for each  $(d'', k) \in (d \rightarrow d'')_i$ 
    for each  $(d', j) \in (d'' \rightarrow d')_k$ 
      if  $(d = d'$  and  $i \neq j)$  or
         $(d > d'$  and  $(d' \rightarrow 0)_j \subseteq (d \rightarrow 0)_i)$ 
           $(d \rightarrow d')_i = (d \rightarrow d')_i \cup (d', j)$ 

```

---

## 4.4 Successive Application of Build, Transpose and Intersection

Any given connectivity  $d \rightarrow d'$  for  $0 \leq d, d' \leq D$  may be computed by a successive application of Algorithms 1–3 in a suitable order. In Algorithm 4, we present the basic logic for a successive and recursive application of the three basic algorithms Build, Transpose and Intersection to compute any given connectivity.

We illustrate this in Figure 5 for computation of the connectivity  $2 \rightarrow 2$ , the incident faces of each face, for a tetrahedral mesh. From the given connectivity  $D \rightarrow 0$ , we first compute the connectivity  $0 \rightarrow D$  by an application of Transpose. This allows us to compute  $D \rightarrow D$  by an application of Intersection. The connectivity  $2 \rightarrow 0$  (and  $D \rightarrow 2$ ) may then be computed by an application of Build. We then apply Transpose to compute  $0 \rightarrow 2$  and finally Intersection to compute  $2 \rightarrow 2$ .

## 4.5 Memory Handling

For each of Algorithms 1–3, memory usage may be conserved by running each algorithm twice; first one round to count the number of incident entities, which allows the static data structures discussed above to be preallocated, and then another round to set the values of the incident




Figure 5: Computing connectivity  $2 \rightarrow 2$  (the faces incident to any given face) by successive application of Transpose, Intersection, Build, Transpose and Intersection.

---

**Algorithm 4**  $\text{Connectivity}(d, d')$ , computing  $d \rightarrow d'$  by application of Algorithms 1–3

---

```

if  $N_d = 0$ 
  Build( $d$ )
if  $N_{d'} = 0$ 
  Build( $d'$ )
if  $d \rightarrow d' \neq \emptyset$ 
  return

if  $d < d'$ 
  Connectivity( $d', d$ )
  Transpose( $d, d'$ )
else
  if  $d = 0$  and  $d' = 0$ 
     $d'' = D$ 
  else
     $d'' = 0$ 
  Connectivity( $d, d''$ )
  Connectivity( $d'', d'$ )
  Intersection( $d, d', d''$ )

```

---

entities. Furthermore, memory usage may be conserved by clearing incidence relations that get computed as byproducts of Algorithms 1–3 when they are no longer needed.

---

## 5 Interfaces

---

In this section, we briefly describe the user interface of the DOLFIN mesh library. We only describe the C++ interface, but note that an (almost) identical Python interface is also available.

### 5.1 Creating a Mesh

A mesh may be created in one of three ways, as illustrated in Figure 6. Either, the mesh is defined by a data file in the DOLFIN XML format<sup>2</sup>, or the mesh is defined vertex by vertex and cell by cell using the DOLFIN mesh editor, or the mesh is defined as one of the DOLFIN built-in meshes. Currently provided built-in meshes include triangular meshes of the unit square and tetrahedral meshes of the unit cube.

```

// Read mesh from file
Mesh mesh0("mesh.xml");

// Build mesh using the mesh editor
Mesh mesh1;
MeshEditor editor;
editor.open(mesh1, "triangle", 2, 2);
editor.initVertices(4);
editor.addVertex(0, 0.0, 0.0);
editor.addVertex(1, 1.0, 0.0);
editor.addVertex(2, 1.0, 1.0);
editor.addVertex(3, 0.0, 1.0);
editor.initCells(2);
editor.addCell(0, 0, 1, 2);
editor.addCell(1, 0, 2, 3);
editor.close();

// Create simple mesh of the unit cube
UnitCube mesh2(16, 16, 16);

```

Figure 6: A DOLFIN mesh may be defined either by an XML data file, or explicitly using the DOLFIN mesh editor, or as a built-in predefined mesh. The last two arguments in the call to `MeshEditor::open()` specify the topological and geometric dimensions of the mesh respectively.

### 5.2 Mesh Iterators

Mesh data may be accessed directly from the mesh, but is most conveniently accessed through the mesh iterator interface. Algorithms operating on a mesh (including Algorithms 1–3) may often be expressed in terms of *iterators*. Mesh iterators can be used to iterate either over the global set of mesh entities of a given topological dimension, or over the locally incident entities of any given mesh entity. Two alternative interfaces are provided; the general interface `MeshEntityIterator` for iteration over entities of some given topological dimension  $d$ , and the spe-

---

<sup>2</sup>A conversion script `dolfin-convert` is provided for conversion from other popular mesh formats (including Gmsh and Medit) to DOLFIN XML format.

cialized mesh iterators `VertexIterator`, `EdgeIterator`, `FaceIterator`, `FacetIterator` and `CellIterator` for iteration over named entities. Iteration over mesh entities may be nested at arbitrary depth and the connectivity (incidence relations) required for any given iteration is automatically computed (at the first occurrence) by the algorithms presented in the previous section.

A `MeshEntityIterator` (`it`) may be dereferenced (`*it`) to create a `MeshEntity`, and any member function `MeshEntity::foo()` may be accessed by `it->foo()`. A `MeshEntityIterator` may thus be thought of as a *pointer* to a `MeshEntity`. Similarly, the named mesh entity iterators may be dereferenced to create the corresponding named mesh entities. Thus, dereferencing a `VertexIterator` gives a `Vertex` which provides an interface to access vertex data. For example, if `it` is a `VertexIterator`, then `it->point()` returns the coordinates of the vertex.

The use of mesh iterators is demonstrated in Figure 7 for iteration over all cells in the mesh and for each cell all its vertices as illustrated in Figure 8. For each cell and each vertex, we print its mesh entity index. We also demonstrate the use of named mesh entity iterators to print the coordinates of each vertex.




Figure 8: Iteration over all vertices of all cells in a mesh. The order of iteration is decided by the definition of the mesh, or alternatively, the UFC ordering convention Alnæs et al. (2007a) if the mesh is ordered. Meshes may be ordered by a call to `Mesh::order()`.

Figure 9 where the same iteration as in Figure 7 is performed without mesh iterators.

```
// Iteration over all vertices of all cells
unsigned int D = mesh.topology().dim();
for (MeshEntityIterator c(mesh, D); !c.end(); ++c)
{
    cout << "cell index = " << c->index() << endl;
    for (MeshEntityIterator v(*c, 0); !v.end(); ++v)
    {
        cout << "vertex index = " << v->index() << endl;
    }
}

// Iteration over all vertices of all cells
for (CellIterator c(mesh); !c.end(); ++c)
{
    cout << "cell index = " << c->index() << endl;
    for (VertexIterator v(*c); !v.end(); ++v)
    {
        cout << "vertex index = " << v->index() << endl;
        cout << "vertex coordinates = " << v->point() << endl;
    }
}
```

Figure 7: Iteration over all vertices of all cells in a mesh, using the general iterator interface `MeshEntityIterator` and the specialized iterators `CellIterator` and `VertexIterator`.

### 5.3 Direct Access to Mesh Data

In addition to the iterator interface, all mesh data may be accessed directly. Thus, one may obtain an array of the vertices of all cells in the mesh directly from the mesh topology, and one may obtain the vertex coordinates of the mesh directly from the mesh geometry. This illustrated in

```
MeshTopology& topology = mesh.topology();
MeshGeometry& geometry = mesh.geometry();
unsigned int D = topology.dim();
MeshConnectivity& connectivity = topology(D, 0);

for (unsigned int c = 0; c < topology.size(D); ++c)
{
    cout << "cell index = " << c << endl;

    unsigned int* vertices = connectivity(c);
    for (unsigned int i = 0; i < connectivity.size(c); ++i)
    {
        unsigned int vertex = vertices[i];

        cout << "vertex index = " << vertex << endl;
        cout << "vertex coordinates = "
            << geometry.point(vertex) << endl;
    }
}
```

Figure 9: Iteration over all vertices of all cells in a mesh and direct access of mesh data corresponding to the iteration of Figure 7 and Figure 8.

### 5.4 Mesh Algorithms

In addition to the computation of mesh connectivity as discussed previously, the DOLFIN mesh library provides a number of other useful mesh algorithms, including boundary extraction, uniform mesh refinement, adaptive mesh refinement (in preparation), mesh smoothing, and reordering of mesh entities.

Figure 10 demonstrates uniform refinement and boundary mesh extraction. When extracting a boundary mesh, it may be desirable to also generate a mapping from the entities of the boundary mesh to the corresponding entities



of the original mesh. This is the case for example when assembling the contribution from boundary integrals during assembly of a linear system arising from a finite element variational formulation of a PDE. One then needs to map each cell of the boundary mesh to the corresponding facet of the original mesh. (Note that the cells of the boundary mesh are facets of the original mesh.) In Figure 11, we demonstrate how to extract a boundary and generate the mapping from the boundary mesh to the original mesh. The mapping is expressed as two `MeshFunctions`, one from the vertices of the boundary mesh to the corresponding vertex indices of the original mesh and one from the cells of the boundary mesh to the corresponding facet indices of the original mesh.

```
// Refine mesh uniformly twice
mesh.refine();
mesh.refine();

// Extract boundary mesh
BoundaryMesh boundary(mesh);

// Refine boundary mesh uniformly
boundary.refine();

// Save boundary mesh to file
File file("boundary.xml");
file << boundary;
```

Figure 10: Uniform refinement, boundary extraction and uniform refinement of the boundary mesh using the DOLFIN mesh library. Note that the extracted boundary mesh is itself a mesh and may thus for example be refined.

```
MeshFunction<unsigned int> vertex_map;
MeshFunction<unsigned int> cell_map;

BoundaryMesh boundary(mesh, vertex_map, cell_map);
```

Figure 11: Extraction of a boundary mesh and generation of a pair of mappings from the vertices of the boundary mesh to the indices of the corresponding vertices of the original mesh and from the cells of the boundary mesh to the indices of the corresponding facets of the original mesh.

## 6 Parallel Considerations

We discuss here how the concepts and data structures discussed above can be used to assemble a global sparse finite element operator (typically matrix) over a mesh distributed over several processors. It is demonstrated below that we may reuse the concepts introduced above to distribute the mesh. In particular, each processor owns a separate piece of the global mesh, which can be stored as a regular `Mesh`. Furthermore, each processor knows which

facets of the local mesh are incident with which facets on other processors and this information can be stored as a pair of `MeshFunctions`.

### 6.1 Simple Distribution of Mesh Data

Let a mesh  $\mathcal{T}$  be given and assume that the mesh has been partitioned into  $n$  disjoint meshes  $\{\mathcal{T}_i\}_{i=0}^{n-1}$  that together cover the computational domain  $\Omega \subset \mathbb{R}^n$ . Such a partition can be computed using for example SCOTCH, see Pellegrini (2004), or Metis, see Karypis and Kumar (1998a,b). On each processor  $p_i$ ,  $i = 0, 2, \dots, n-1$ , we store its part of the global mesh as a regular mesh and in addition two mesh functions  $\mathcal{S}_i$  and  $\mathcal{F}_i$  over the facets of  $\mathcal{T}_i$ .

We thus propose to store a distributed mesh  $\mathcal{T}$  on a set of processors as the set of tuples  $\{(\mathcal{T}_i, \mathcal{S}_i, \mathcal{F}_i)\}_{i=0}^{n-1}$  where one tuple  $(\mathcal{T}_i, \mathcal{S}_i, \mathcal{F}_i)$  is stored on each processor  $p_i$  for  $i = 0, 1, \dots, n-1$ .

The mesh function  $\mathcal{S}_i$  maps each facet  $f$  to an integer  $j = \mathcal{S}_i(f)$  which indicates which (other) subdomain/mesh  $\mathcal{T}_j$  that the facet  $f$  is (physically) incident with,

$$\mathcal{S}_i : (D-1, [0, 1, \dots, N_{D-1}^i - 1]) \rightarrow [0, 1, \dots, n-1]. \quad (5)$$

Here,  $(D-1, [0, 1, \dots, N_{D-1}^i - 1])$  indicates that the domain of  $\mathcal{S}_i$  is the set of tuples  $\{(D-1, k)\}$  where  $0 \leq k \leq N_{D-1}^i - 1$  and  $N_{D-1}^i$  is the number of facets of  $\mathcal{T}_i$ . Thus, if  $j = \mathcal{S}_i(f)$  for some  $j \neq i$ , then the facet  $f$  is shared with the mesh  $\mathcal{T}_j$ . If the facet  $f$  is only incident with  $\mathcal{T}_i$  itself, then we set  $\mathcal{S}_i(f) = i$ .

The mesh function  $\mathcal{F}_i$  maps each facet entity  $f$  to an integer  $\mathcal{F}_i(f)$  which indicates which facet  $f' = (D-1, \mathcal{F}_i(f))$  of  $\mathcal{T}_j$  for  $j = \mathcal{S}_i(f)$  that the facet  $f$  is incident (identical) to. If  $\mathcal{S}_i(f) = i$ , then  $f$  is not shared with another mesh and we set  $\mathcal{F}_i(f) = 0$ . We illustrate the meaning of the two mesh functions  $\mathcal{S}_i$  and  $\mathcal{F}_i$  in Figure 12.

### 6.2 Parallel Assembly

The standard algorithm for computing a global sparse operator (tensor) from a finite element variational form is known as *assembly*, see Zienkiewicz et al. (1967); Hughes (1987); Langtangen (1999). By this algorithm, the global sparse operator may be computed by assembling (summing) the contributions from the local entities of a finite element mesh. On each cell of the mesh, one computes a small *cell tensor* (often referred to as the “element stiffness matrix”) and add the entries of that tensor to a global sparse tensor (often referred to as the “global stiffness matrix”). We shall not discuss the assembly algorithm in detail here and refer instead to Alnæs et al. (2008); Logg (2007); Alnæs et al. (2007b), but note that to add the entries from the local cell tensor to the global sparse tensor, we need to compute a so-called *local-to-global* mapping on each cell. This maps the local degrees of freedom on a cell (numbering the rows and columns of the cell tensor) to global degrees of freedom (numbering the rows and columns of the global tensor).




Figure 12: A global mesh  $\mathcal{T}$  partitioned into  $n = 3$  meshes  $\{\mathcal{T}_i\}_{i=0}^{n-1}$ . The mesh functions  $\mathcal{S}_i$  and  $\mathcal{F}_i$  indicate which facets of  $\mathcal{T}_i$  are shared with other meshes. In this example, we have  $\mathcal{S}_0((1,3)) = \mathcal{S}_0((1,4)) = 1$ , indicating that facets  $(1,3)$  and  $(1,4)$  in  $\mathcal{T}_0$  are shared with  $\mathcal{T}_1$ . Furthermore, we may evaluate  $\mathcal{F}_0$  at these two facets to find that the facet  $(1,3)$  in  $\mathcal{T}_0$  is incident to facet  $(1,4) = \mathcal{F}_0((1,3))$  in  $\mathcal{T}_1$  and facet  $(1,4)$  in  $\mathcal{T}_0$  is incident to facet  $(1,5) = \mathcal{F}_0((1,4))$  in  $\mathcal{T}_1$ .

The assembly algorithm may be trivially parallelized by letting each processor  $p_i$  compute and insert the cell tensors on the local mesh  $\mathcal{T}_i$  into the global tensor. When the global tensor is a sparse matrix, linear algebra libraries like PETSc, see Balay et al. (2004, 1997), may be used to store the global tensor in parallel. PETSc handles the communication of matrix data between processors and we need only make sure that each processor knows how to insert entries into the global sparse matrix according to the local-to-global mapping of the global mesh. We demonstrate below that on each processor  $p_i$ , we may (with a small amount of communication with neighboring processors) compute the part of the local-to-global mapping of the global mesh  $\mathcal{T}$  relevant to each local mesh  $\mathcal{T}_i$  in parallel on each processor  $p_i$ , which thus allows us to assemble the global sparse matrix in parallel.

### 6.3 Mapping Degrees of Freedom in Parallel

In Algorithm 5, we describe how the mapping of degrees of freedom may be computed on a distributed mesh  $\{(\mathcal{T}_i, \mathcal{S}_i, \mathcal{F}_i)\}_{i=0}^{n-1}$ .

To express this algorithm in compact form, we need to introduce some further notation. For each cell  $c$  in a local finite element mesh  $\mathcal{T}_i$ , we assume that we can compute a local-to-global mapping  $\iota_c^i$ , which maps each local degree of freedom on the cell  $c$  to a global degree of freedom (for a numbering scheme valid on the local mesh  $\mathcal{T}_i$ ). For example, when computing with standard piecewise linear finite elements on triangles, the local-to-global mapping  $\iota_c^i$  may map a local vertex number 0, 1 or 2 on  $c$  to the corresponding global number vertex number on the mesh  $\mathcal{T}_i$ . Thus, the domain of  $\iota_c^i$  is here  $\{0, 1, 2\}$  and the range is  $[0, N_0^i - 1]$ , where  $N_0^i$  is the number of vertices of the mesh  $\mathcal{T}_i$ . We em-

phasize that the local-to-global mapping  $\iota_c^i$  is not aware of the global mesh  $\mathcal{T}$  of which the local mesh  $\mathcal{T}_i$  is a part. Instead, it is the task of Algorithm 5 to compute (in parallel) a local-to-global mapping valid on the global mesh  $\mathcal{T}$  from a given local-to-global mapping on each local mesh  $\mathcal{T}_i$ .

We let  $\mathcal{M}_i$  denote the parallel local-to-global mapping to be computed on each part  $\mathcal{T}_i$  of the global mesh. For ease of notation, we express  $\mathcal{M}_i$  as a set of tuples  $\mathcal{M}_i = \{((c, i), I)\}$ , where  $c$  is a cell (number),  $i$  is the local number of a degree of freedom on  $c$  and  $I$  is the corresponding global number. The mapping  $\mathcal{M}_i$  should be thought of as a function that maps a cell and a local degree of freedom  $(c, i)$  to the corresponding global degree of freedom  $I$ . In the case of standard piecewise linear elements on triangles, the domain of  $\mathcal{M}_i$  is  $[0, N_D^i - 1] \times \{0, 1, 2\}$ , where  $N_D^i$  is the number of cells of the mesh  $\mathcal{T}_i$ , and the range of  $\mathcal{M}_i$  is  $[0, N_0 - 1]$ , where  $N_0$  is the total number of vertices of the partitioned global mesh  $\mathcal{T}$ . We note here that the mapping  $\mathcal{M}_i$  may be stored as a fixed-size array.

To compute the parallel local-to-global mapping  $\mathcal{M}_i$  on each processor, we need to iterate over the entities of the mesh  $\mathcal{T}_i$  and renumber the degrees of freedom. To do this, we introduce an auxiliary temporary mapping  $\mathcal{N}_i$  on each processor that maps degrees of freedom as given by the local-to-global mapping  $\iota_c^i$  on each cell  $c$  (which is only aware of how to map degrees of freedom internally on  $\mathcal{T}_i$ ) to degrees of freedom as given by the parallel local-to-global mapping  $\mathcal{M}_i$  (which is aware of how to map degrees of freedom globally on the distributed mesh). In the case of standard piecewise linear elements on triangles, the domain of  $\mathcal{N}_i$  is  $[0, N_0^i - 1]$  and the range of  $\mathcal{N}_i$  is  $[0, N_0 - 1]$ . Just as for  $\mathcal{M}_i$ , we may think of  $\mathcal{N}_i$  as a function but write it as a set of tuples (pairs) in Algorithm 5 for ease of notation. In a C++ implementation,  $\mathcal{N}_i$  may be stored in the form of an STL map (`std::map<unsigned int, unsigned int>`).

Finally, we let  $\overline{N}_i$  denote the number of degrees of freedom on  $\mathcal{T}_i$  not shared with a mesh  $\mathcal{T}_j$  for  $j < i$ . This number can be computed on each mesh  $\mathcal{T}_i$  from the mesh function  $\mathcal{S}_i$ .

In Algorithm 5, we first let each processor  $p_i$  compute  $\overline{N}_i$  (in parallel). These numbers are then communicated successively from  $p_{i-1}$  to  $p_i$  to compute an offset for the numbering of degrees of freedom on each mesh  $\mathcal{T}_i$ . We then let each processor  $p_i$  number the degrees of freedom (in parallel) on cells which are incident with the boundary of  $\mathcal{T}_i$  and are *shared with a mesh  $\mathcal{T}_j$  for  $j > i$* . After the degrees of freedom on mesh boundaries have been numbered on each processor, those numbering schemes are communicated successively from  $p_i$  to  $p_j$  for all  $i < j$  such that  $\mathcal{T}_i$  and  $\mathcal{T}_j$  share degrees of freedom on a common facet. Finally, each processor  $p_i$  may number the  $\overline{N}_i$  “internal degrees of freedom” on  $\mathcal{T}_i$ .

The key point of this algorithm is to always let  $p_i$  number degrees of freedom common with  $p_j$  for  $i < j$ . This numbering of shared degrees of freedom is then communicated over shared facets, from  $f'$  to  $f$  in stage 2 of Algorithm 5. Since it is known a priori which facets are shared between

two meshes  $\mathcal{T}_i$  and  $\mathcal{T}_j$ , one may communicate the common numbering for all shared degrees of freedom from  $p_i$  to  $p_j$  in one batch.

Note that it is important that the communication of shared degrees of freedom in stage 2 of Algorithm 5 is performed sequentially, starting with processor  $p_1$  receiving the common numbering from  $p_0$ , then  $p_2$  receiving the common numbering from  $p_0$  and/or  $p_1$  etc. This guarantees that common degrees of freedom are communicated from  $p_i$  to all  $p_j$  with  $i < j$  such that  $\mathcal{T}_i$  and  $\mathcal{T}_j$  share common degrees of freedom, even if  $\mathcal{T}_i$  and  $\mathcal{T}_j$  don't share a common facet. For this to work, we make the assumption that if any two meshes  $\mathcal{T}_i$  and  $\mathcal{T}_j$  share a common degree of freedom on the two cells  $c \in \mathcal{T}_i$  and  $c' \in \mathcal{T}_j$ , then each of  $\mathcal{T}_i$  and  $\mathcal{T}_j$  must share that degree of freedom with some other mesh  $\mathcal{T}_{i'}$  or  $\mathcal{T}_{j'}$  respectively. In Figure 13, we illustrate this assumption (for  $i' = j'$ ) and give an example of a partition for which Algorithm 5 fails to correctly number all shared degrees of freedom. It is a mild assumption to disallow such partitions (and meshes).

Algorithm 5 is currently not implemented in DOLFIN. Instead, a simple (but suboptimal) strategy where the computational mesh is broadcast to all processors has been implemented. Each processor owns a copy of the entire mesh and knows which part of that mesh to assemble. For a further discussion of the current implementation of parallel assembly in DOLFIN (available with DOLFIN 0.7.2), see Vikstrøm (2008).




Figure 13: Two partitions of a mesh. In the partition on the left, two of the meshes share only a common vertex and may thus share a single degree of freedom at that vertex. The communication of a common numbering of that degree of freedom is propagated over the facets incident with the common vertex to the neighboring mesh. In the partition on the right, it is not possible to propagate the numbering over facets (since there are no shared facets) and so Algorithm 5 will fail to compute a correct numbering scheme for this partition.

---

## 7 Benchmark Results

---

In this section, we present a series of benchmarks to illustrate the efficiency of the mesh representation and its implementation. The new mesh library (which is available

---

**Algorithm 5**  $\{\mathcal{M}_i\} = \text{ComputeMapping}(\{(\mathcal{T}_i, \mathcal{S}_i, \mathcal{F}_i)\})$ , computing the local-to-global mapping in parallel for a mesh distributed over  $n$  processors  $p_i$ ,  $i = 0, 1, \dots, n-1$ .

---

— Stage 0: *Compute offsets*

**on each** processor  $p_i$

    compute  $\overline{N}_i$

**on** processor  $p_0$

    offset<sub>0</sub> = 0

**for**  $i = 1, 2, \dots, n-1$

**on** processor  $p_i$

        Receive (offset <sub>$i-1$</sub> ,  $\overline{N}_{i-1}$ ) from  $p_{i-1}$

        offset <sub>$i$</sub>  = offset <sub>$i-1$</sub>  +  $\overline{N}_{i-1}$

— Stage 1: *Compute mapping on shared facets*

**on each** processor  $p_i$

$\mathcal{M}_i = \emptyset$

$\mathcal{N}_i = \emptyset$

$k_i = \overline{N}_i$

**for each** facet  $f \in \mathcal{T}_i$

$j = \mathcal{S}_i(f)$

**if**  $j > i$

            Let  $c \in \mathcal{T}_i$  be the cell incident with  $f$

**for each** local degree of freedom  $l$  on  $c$

**if**  $(\iota_c^i(l), L) \in \mathcal{N}_i$  for some  $L$

$\mathcal{M}_i = \mathcal{M}_i \cup ((c, l), L)$

**else**

$\mathcal{M}_i = \mathcal{M}_i \cup ((c, l), k_i)$

$\mathcal{N}_i = \mathcal{N}_i \cup (\iota_c^i(l), k_i)$

$k_i = k_i + 1$

— Stage 2: *Communicate mapping on shared facets*

**for**  $j = 1, 2, \dots, n-1$

**for each** facet  $f \in \mathcal{T}_j$

$i = \mathcal{S}_j(f)$

**if**  $i < j$

$f' = \mathcal{F}_j(f)$

            Receive degrees of freedom on  $f'$  from  $p_i$

            Update  $\mathcal{N}_j$  for shared degrees of freedom

— Stage 3: *Compute mapping for interior degrees of freedom*

**on each** processor  $p_i$

**for each** cell  $c \in \mathcal{T}_i$

**for each** local degree of freedom  $l$  on  $c$

**if**  $(\iota_c^i(l), L) \in \mathcal{N}_i$  for some  $L$

$\mathcal{M}_i = \mathcal{M}_i \cup ((c, l), L)$

**else**

$\mathcal{M}_i = \mathcal{M}_i \cup ((c, l), k_i)$

$\mathcal{N}_i = \mathcal{N}_i \cup (\iota_c^i(l), k_i)$

$k_i = k_i + 1$

---

as part of DOLFIN since version 0.6.3) is compared to the old DOLFIN mesh library which is a fairly efficient C++ implementation, but which suffers from object-oriented overhead; all mesh entities are there stored as arrays of objects, which store their data locally in each object (including mesh incidence relations and vertex coordinates). The benchmark results were obtained on a 2.66 GHz 64-

bit Intel processor (Q6700) running Ubuntu GNU/Linux for DOLFIN version 0.6.2-1 and DOLFIN version 0.7.1 respectively.

The five test cases that are examined are the following: (i) CPU time and (ii) memory usage for creation of a uniform tetrahedral mesh of the unit cube, (iii) CPU time for iteration over all vertices of the mesh, (iv) CPU time for accessing the coordinates of all vertices of the mesh, and (v) uniform refinement of the mesh.

In summary, the speedup was in all cases a factor 10–1000 and memory usage was reduced by a factor of 50 (comparing slopes of lines in Figure 15). The speedup and decreased memory usage is the result of more efficient algorithms and data structures, where all data is stored in large static arrays and objects are only provided as part of the interface for simple access to the underlying data representation, not to store data themselves. Another contributing factor is that the old DOLFIN mesh library precomputes certain connectivities (including the edges and faces of each cell) at startup, whereas this computation is carried out only when requested in the new DOLFIN mesh library, either as part of the iterator interface or by an explicit call to `Mesh::init()`. We also note that from Figure 15, it is evident that DOLFIN can be further improved in terms of its minimal memory footprint.

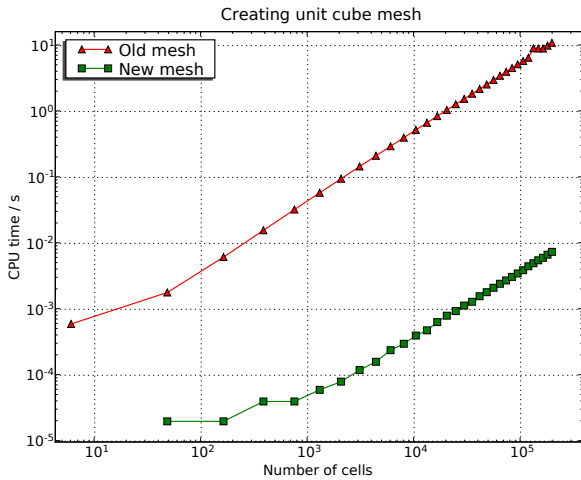


Figure 14: Benchmarking the CPU time for creation of a uniform tetrahedral mesh of the unit cube for the new mesh library vs. the old DOLFIN mesh library.

## 8 Conclusions

We have presented a simple yet general and efficient representation of computational meshes and demonstrated a straightforward implementation of this representation as a set of C++ classes that correspond to the basic concepts of the mesh representation. The implementation is available freely as part of DOLFIN Hoffman et al. (2006).

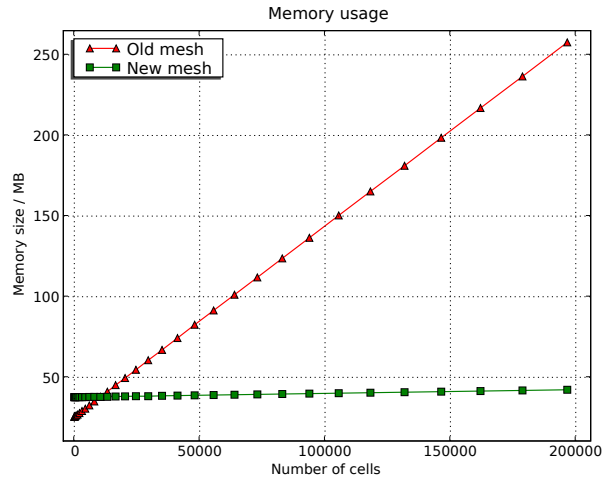


Figure 15: Benchmarking the memory usage for creation of a uniform tetrahedral mesh of the unit cube for the new mesh library vs. the old DOLFIN mesh library.

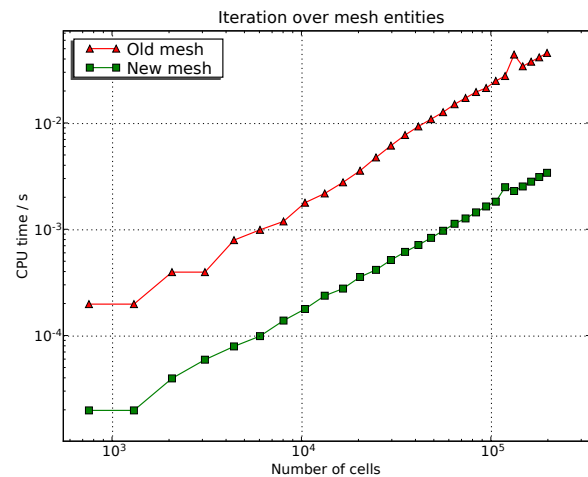


Figure 16: Benchmarking the CPU time for iteration over the vertices of each cell for the new mesh library vs. the old DOLFIN mesh library.

---

## REFERENCES

---

- Alnæs, M., Langtangen, H. P., A. Logg, Mardal, K.-A., and Skavhaug, O. (2007a). UFC. URL: <http://www.fenics.org/ufc/>.
- Alnæs, M., Langtangen, H. P., A. Logg, Mardal, K.-A., and Skavhaug, O. (2007b). *UFC Specification and User Manual*. URL: <http://www.fenics.org/ufc/>.
- Alnæs, M. S., Langtangen, H. P., Logg, A., Mardal, K.-A., and Skavhaug, O. (2008). Unified framework for finite element assembly. *Submitted to International Journal of Computational Science and Engineering*.




Figure 17: Benchmarking the CPU time for iteration over the coordinates of each vertex for the new mesh library vs. the old mesh library vs. direct access of coordinate arrays in the new DOLFIN mesh library.

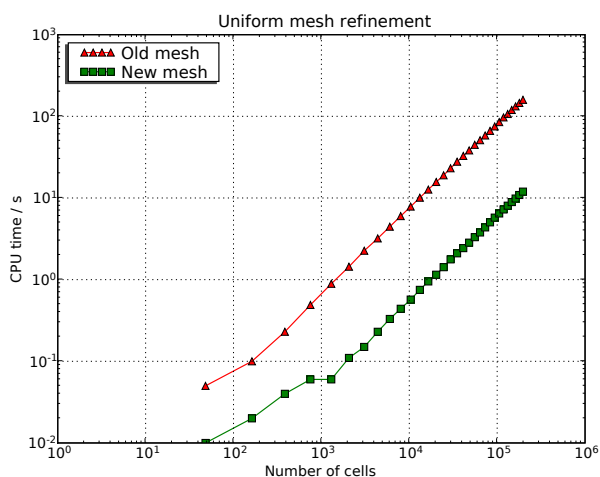


Figure 18: Benchmarking the CPU time for uniform refinement of the unit cube for the new mesh library vs. the old DOLFIN mesh library.

Balay, S., Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H. (2004). PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory.

Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F. (1997). Efficient management of parallelism in object oriented numerical software libraries. In Arge, E., Bruaset, A. M., and Langtangen, H. P., editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press.

Berti, G. (2002). Generic programming for mesh algorithms: Towards universally usable geometric components. In *Proceedings of the Fifth World Congress on Computational Mechanics, Vienna University of Technology July*.

Berti, G. (2006). GrAL - the grid algorithms library. *Future Generation Computer Systems*, 22.

Dupont, T., Hoffman, J., Johnson, C., Kirby, R. C., Larson, M. G., Logg, A., and Scott, L. R. (2003). The FEniCS project. Technical Report 2003-21, Chalmers Finite Element Center Preprint Series.

FEniCS (2008). FEniCS project. URL: <http://www.fenics.org/>.

Hoffman, J., Jansson, J., Logg, A., and Wells, G. N. (2006). DOLFIN. <http://www.fenics.org/dolfin/>.

Hughes, T. J. R. (1987). *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice-Hall.

Karpeev, D. A. and Knepley, M. G. (2005). Flexible representation of computational meshes. *Submitted to ACM Trans. Math. Softw.*

Karypis, G. and Kumar, V. (1998a). A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20(1):359–392.

Karypis, G. and Kumar, V. (1998b). *MeTis: Unstructured Graph Partitioning and Sparse Matrix Ordering System, Version 4.0*.

Langtangen, H. P. (1999). *Computational Partial Differential Equations – Numerical Methods and Diffpack Programming*. Lecture Notes in Computational Science and Engineering. Springer.

Logg, A. (2007). Automating the finite element method. *Arch. Comput. Methods Eng.*, 14:93–138.

Pellegrini, F. (2004). Native mesh ordering with SCOTCH 4.0. In *Proceedings of VECPAR'04, Valencia, Spain*. soumis.

Vikström, G. M. (2008). Parallelization strategies for DOLFIN.

Zienkiewicz, O. C., Taylor, R. L., and Zhu, J. Z. (2005, first published in 1967). *The Finite Element Method — Its Basis and Fundamentals, 6th edition*. Elsevier.