

Binarized octree generation for Cartesian adaptive mesh refinement around immersed geometries

Jaber J. Hasbestan¹, Inanc Senocak²

*Department of Mechanical Engineering and Materials Science,
University of Pittsburgh, Pittsburgh, PA 15261, USA*

Abstract

We revisit the generation of balanced octrees for adaptive mesh refinement (AMR) of Cartesian domains with immersed complex geometries. In a recent short note [Hasbestan and Senocak, J. Comput. Phys. vol. 351:473-477 (2017)], we showed that the data-locality of the Z-order curve in hashed linear octree generation methods may not be perfect because of potential collisions in the hash table. Building on that observation, we propose a binarized octree generation method that complies with the Z-order curve exactly. Similar to a hashed linear octree generation method, we use Morton encoding to index the nodes of an octree, but use a red-black tree in place of the hash table. Red-black tree is a special kind of a binary tree, which we use for insertion and deletion of elements during mesh adaptation. By strictly working with the bitwise representation of the octree, we remove computer hardware limitations on the depth of adaptation on a single processor. Additionally, we introduce a geometry encoding technique for rapidly tagging the solid geometry for refinement. Our results for several geometries with different levels of adaptations show that the binarized octree generation outperforms the linear octree generation in terms of runtime performance at the expense of only a slight increase in memory usage. We provide the current AMR capability as open-source software.

Keywords: AMR, Hash table, Octree, Red-black tree, Z-order curve

1. Introduction

Mesh adaptation is a powerful technique in the solution of multiphysics problems governed by coupled partial differential equations (PDEs). By refining the spatial resolution where needed, researchers can realize substantial savings in computer memory and solution turnaround time. Unlike the solution algorithm, mesh adaptation is not governed by a set of governing equations. Therefore, different approaches and strategies have been proposed for mesh adaptation [1].

¹Postdoctoral Research Associate, jaber@pitt.edu

²Associate Professor, senocak@pitt.edu

In this study our focus is on the so-called structured adaptive mesh refinement (AMR) method. Berger and Oliger [2] pioneered the use of patch-based adaptation for the solution of hyperbolic PDEs with refinement in both space and time. Berger and Colella [3] further improved the AMR algorithm and used an estimate of the local truncation error to guide the refinement decisions. Although these pioneering works were restricted to an explicit finite-difference method, they are one of the early works that introduced efficient use of data-structures in scientific computing and the strategy developed in these works is often referred to as *Berger-Oliger-Colella adaptive refinement*. This unique patch-based strategy has been adopted in Boxlib [4], PARAMESH [5], AMREx [6], Chombo [7], ENZO [8], SAMRAI [9, 10] and AMRClaw [11] packages. In a patch-based AMR, a hierarchy of patches are nested within a coarse level mesh that occupies the entire domain. A refinement ratio is assigned between successive patches. During refinement, patches are tagged based on a specified criterion for refinement. Ghost-cells are used at patch boundaries to compute the solution across the patches.

Implementation of an AMR algorithm is major software development effort. It is common to refer to the end product as an AMR infrastructure. In building this software infrastructure, the design goal is to introduce an abstraction for the meshing part such that developers of the solution algorithm need not know the details of the computer implementation. For static-mesh problems, the efficiency of the AMR infrastructure may not be critical. But for large-scale problems in which the refinement criteria is dynamic, time spent on meshing, while the simulation is evolving, can be prohibitively expensive. Additionally, the memory hierarchy on modern computing hardware favors computer implementations that exploit data-locality and use less memory. It is worth mentioning that these performance issues related to adaptive meshing have been the focus of several studies in the computer graphics field, driven by real-time applications [12]. There are a number of open-source packages that provide AMR capabilities for meshing and solving PDEs. Dubey et al. [13] survey a representative set of block-structured AMR infrastructures.

An alternative strategy to the patch-based AMR approach is a tree-based strategy. Octree is a hierarchical data structure that has been the subject of many studies [14, 15]. The tree structure has also been adopted in the solution of partial differential equations (PDEs) with adaptive resolutions. In its classical representation, an octree keeps track of parent-child relations with pointers. In the three-dimensional case, eight pointers are needed to store the neighbors.

Sundar et al. [16] developed parallel algorithms to generate balanced linear octrees. A linear octree is a representation of an octree in which a linear array is used instead of a tree data structure to store the leaf nodes [17, 18]. It is common to use space-filling curves to represent linear octrees. Sundar et al. used Morton encoding as one of the many ways to generate a space-filling curve because of its ease of computer implementation. Their method addressed the so-called *ripple-effect* associated with balancing the octree.

Burstedde et al. [19] pursued the linear octree representation approach described in Sundar et al. [16] and developed `p4est`, which is an AMR package

for multiple connected octrees. Authors refer to these connected octrees as *forest-of-octrees*. An encoding scheme was introduced to enable inter-octree connectivity. Octrees and octants were numbered using a space-filling Z-order curve, which was effective in parallel load balancing. `p4est` has been shown scale across many processors. In one exercise, a mesh with 5.13×10^{11} elements was created using 220,320 processes. Isaac et al. [20] introduced recursive algorithms for this forest-of-octrees approach to traverse a linear octree and further improve the parallel scalability of `p4est` aimed at improving the 2:1 balance algorithm.

Another approach is the building-cube method (BCM) [21, 22], which is a combination of explicit tree and unstructured grid methodology. In BCM, each cube has its own equally-spaced Cartesian grid. All cubes have the same number of Cartesian mesh. Therefore, the resolution is determined by the level of adaptation of the cube.

Having a balanced tree is important in numerical solution of PDEs as it enables one to gradually traverse from low-resolution to high-resolution without significantly degrading numerical accuracy, especially for simulation of multi-scale problems such as turbulent flows. A balanced-tree is also needed to enable a generic implementation of numerical schemes for the solution of PDEs in a code because coarse to fine blocks obeys the same rule. The common approach to enforce a balanced tree is to confine the maximum level of the element at each face to 2^{d-1} where d is the dimension of the computational space (i.e. 2 or 3). The difficulty of the 2:1 balancing algorithm emanates from the fact that the domain of influence of any element under consideration for refinement might extend beyond the immediate neighboring elements. This issue has been coined as the *ripple effect* in the literature [16]. Different methods have been proposed to improve the efficiency of enforcing the 2:1 balance condition.

Tu and Ohallaron [23] demonstrate 2:1 balanced refinement of octrees. Their solution consists of two steps; balance by parts and prioritize the *ripple effect*. The key idea in their approach is to bulk load most of the data on memory and enforce the balance locally. They load a small region to a pointer-based (sub)octree in memory and resolve the 2:1 constraint and the ripple effect without worrying about octants outside of the region. This is due to the fact that most impact caused by a tiny octant tagged to be refined is localized in a small region. Their balance-by-parts algorithm is described as follows. Imagine a moving window inside the 3D domain, at any given moment, the octants that reside in this window are retrieved from disk and cached in a (sub)octree. Then, they adjust the data structure to enforce 2-to-1 constraint in memory, the content on disk is updated accordingly.

Ripple effect is even more challenging in distributed systems due to the complicated communication pattern. Sundar et al. [16] proposed a bottom-up approach to enforce the balance condition. This approach eliminates the need for synchronization in parallel implementation. Isaac et al. [20] improved the balance algorithm in the `p4est` package. In their approach, they first reduce the input set to a compressed equivalent and iteratively add a sparse set of octants to that list to enforce the balance.

Aside from solving the PDEs, octrees are extensively used in computer graphics. Unlike in the solution of PDEs, balanced trees are not required in computer graphics applications, and the desired resolution can be much lower than resolutions needed for turbulent flow simulations. Driven by interactivity requirements of computer games and ray tracing applications, the implementation of octree on central processing units (CPUs) or graphics processing units (GPUs) can be highly optimized. Morton codes are used in building octrees, mapping multi-dimensional data to one dimension and thereby enhancing data-locality. Use of a hash table with Morton encoding is common in generating the octrees. Meagher [14] utilized octree data structure for modeling the three-dimensional complex geometries including geometries with holes and disjointed parts and sculptured surfaces. Frisken and Perry [24] demonstrate efficient methods for tree structure traversals.

Lewiner et al. [12] developed a fast pointerless method to generate dual octrees. They proposed strategies for dual generation of static or dynamic pointerless octrees, where memory usage is lowered by several factors compared to the usual recursive generation methodologies. Note that pointerless representation is another terminology for hashed linear octrees. However, their study is not related to solving PDEs and hence, the generated mesh is not necessarily balanced. Lewiner et al. provide a good review on different approaches to octree generation in the same work.

Schwarz and Seidel [25] describe octree implementation on GPUs. Baert et al. [26] uses bit-interleaving in constructing the Morton code to map every voxel center to a one-dimensional space to generate out-of-core octrees for a Cartesian mesh, the method presented there is not directly applicable to AMR, as they use integer indices in the mapping. The bit-interleaving introduced in that study uses fixed point arithmetic to store the center of the cube. Bit-interleaving is also introduced by Tero Karras [27]. The problem with these approaches for solving PDEs is that one needs to always scale the geometry to unit length and extracting the connectivity is not trivial.

Different versions of space-filling curves have been used for partitioning the octrees [28, 29, 30, 31]. Z-order curve, also referred to as Morton code [32], is often preferred because of its ease of computer implementation [16]. An important feature of using Morton code is that it ensures that the elements with closer Morton codes are stored close to each other in the computer memory. This feature makes the search for a given key efficient, since one only needs to look for the elements close to the current key.

The method used to search for a key is also critical for performance. In a recent short note [33], we have shown that a red-black tree helps preserve the Z-order curve exactly, whereas the Z-order can be broken because of collisions in a hashed linear octree. A red-black tree also supports deep-level mesh adaptation for extreme-scale problems because the method is based on bitwise comparisons only.

Building on those useful features of the red-black tree data structure, we present a fast, low-storage AMR algorithm that complies with the Z-order curve strictly. Our presentation includes the algorithms for 2:1 balance, geometry

encoding for immersion of complex geometries, and the tricks and trades of computer implementation that exploits bitwise representation heavily. In our approach, we do not store the neighborhood connectivity explicitly. To the best of our knowledge, our open-source software, which we call `rebl-AMR`, is the first Cartesian adaptive mesh generation package that make use of the fundamental red-black tree data structure at its core. We present numerous examples to demonstrate the performance of red-black tree version and compare its computational performance to a hashed linear octree. For completeness, we compare the performance results from both methods against a graph-based unstructured mesh approach as well. The present `rebl-AMR` [34] package is available as open-source software.

2. Octree representations

An octree is constructed by isotropic refinement where a rectangular volume is subdivided in half in x , y and z directions. Refinement of each element generates eight more elements in three dimensions. We refer to these elements as *siblings*. Complicated geometries can be represented using octrees by recursively subdividing the elements. Different types of data structures have been used to store the tree data. The selection of the data structures is a crucial step in octree generation algorithms because it will impact the execution time as well as memory usage. For dynamic adaptive mesh generation problems where the geometry deforms or moves in time, it is desirable to have constant-in-time insertion and removal of mesh elements.

One of the early approaches in octree generation is the use of parent-child relation in creating a so-called explicit octree, where use of pointers to children is standard. This approach is the classical way of storing the tree and uses eight pointers for octree representation [35]. A sample four-level tree is presented in Fig. 4. The number of pointers needed to represent a large adaptively refined mesh would be unnecessarily large as well and would also complicate the computer implementation. Pointerless implicit trees have been proposed as an alternative to the explicit tree representation [17, 18].

We present three methods to generate octrees. The first method is the hashed linear octree generation. Although this method is very efficient, there is no guarantee to preserve the Z-order curve once it is used in conjunction with a hash table because of collision, which we have demonstrated recently [33]. The second method is unique to the present work, which we call it the *binarized* octree generation. Similar to a hashed linear octree generation method, we represent octree nodes on a Z-order curve and use a red-black tree for element search, insertion and deletion operations. A red-black is a special kind of binary-tree and bitwise representation is used to implement the Z-order curve. Because we use a binary tree data structure to store the leaves of the octree instead of an array, we refer to this approach as *binarized octree generation*. Compared to the classical approach to generate an octree, both the hashed linear octree and the binarized octree use much less memory.

In the third approach, we generate an octree as an unstructured mesh. The connectivity is generated for every element using a graph-based approach. This method updates the connectivity directly without the need to actually store the tree information explicitly. The unstructured mesh approach is provided for completeness. We describe that approach in Hasbestan and Senocak [36] and it is not the main focus of the current work.

Z-order curve, also referred to as the Morton code, is common to both the linear and binarized octree generation methods. Therefore, we first explain the construction of Morton codes and how we manipulate them for neighborhood connectivity information.

2.1. Morton code generation

Morton encoding is a mapping from a multi-dimensional space to one dimension [32]. When generating a Morton code, first, a bit code is constructed for every node. This node is then converted to an integer, if needed. The nodes, once laid out, follow a Z-order curve, which enhances data-locality.

Figure 1 presents the Morton ordering for a mesh in three dimensions, which is refined one level only. Note that we colored the four lower octants on their bottom surface while coloring the top four octants on their top surfaces to avoid visual clutter. The bit representation on each surface is valid for the individual octant. The Morton code for each octant is constructed by specifying its relative location index in x, y, z directions in the Cartesian coordinates, respectively.




Figure 1: Morton code construction at a single level based on element location in x, y, z directions

Let us now demonstrate refinement. At each level, the Morton code corresponds to three bits in three-dimensional space and two bits in two-dimensional space. The Morton code after refinement is generated by concatenating every bits at every level, starting from the left with the top level node in the tree and placing each node's bit representation to the right as we traverse to the bottom of the tree. A two-dimensional mesh with two levels of adaptation is

presented in Fig. 2 to illustrate Morton code generation with refinement. Here, we considered a two-dimensional case for ease of presentation. The final code is generated by concatenating two bits per level of adaptation. The length of the codes are extended to 4 bits by zero padding to be consistent.




Figure 2: Morton code representation of a two-dimensional mesh with two levels of adaptation.

An octree for a three-dimensional mesh with three levels of adaptation is shown in Fig. 3. Here, the procedure is illustrated with coordinated colors for node 23. The Morton code for node number 23 is 001,000,101 as shown in Fig. 4. Note that the change of color in Fig. 4 coordinates with Fig. 3 to highlight placement of bits from left to right.

Morton code enables us to keep the elements that have closer keys allocated close to each other in the hash table. It is important to mention that the cache performance also improves due to this locality. There are several methods for Morton encoding [27, 12].

The procedure described above limits the number of levels for adaptation to 10 for natively 32-bit systems and to 21 for natively 64-bit systems. Our statement is valid if the AMR algorithm at hand has to convert the Morton code to an integer number because it would produce an overflow.

We should note that in the hashed linear octree approach, any hash function that does not use the entire bits and only selects some portion of the Morton code would be able to attain adaptation levels deeper than the aforementioned values at the expense of performance due to collisions in the hash table. As we have shown in Hasbestan and Senocak [33] Z-order curve is no longer retained once collisions occur. Lewiner et al. [12] discuss several approaches pertinent to the use of hash table with Morton codes.

As we discuss later in this paper, we can overcome this limitation on depth of adaptation without a performance hit by using the red-black tree data structure in the AMR algorithm. C++ `bitset` class enables defining the number of required bits beyond what is natively supported by the hardware, and the red-black tree operates on bit comparison and therefore there is no need for




Figure 3: Morton code construction procedure illustrated for the node 23, The numbers in the parentheses represent the node index as an integer number

001,000,101

Figure 4: Morton code constructed for node 23 after three levels of refinement

converting the Morton code to an integer.

2.2. Hashed linear octree generation

Data-locality is key to superior memory performance. As we mentioned earlier, one option to enhance data-locality is to store data using space-filling curves. The use of Z-order curve [32] is popular because of its easy computer implementation. Use of hash tables is also common in pointerless linear octree generation. A hash table potentially provide $O(1)$ access to its elements owing to the use of array as the underlying data-structure. However, $O(1)$ access is only guaranteed as long as there are no collisions in the hash table. In practice, this requirement is difficult to fulfill and collisions do happen when the hash function generates the same index for different keys. Typically, a remainder operator is used as the hash function to fit the keys inside a given array.

The state-of-the-art in linear octree generation is to use a hash table with Morton codes. This choice combines the efficiency of a hash table with data-locality offered by the Z-order curve to further improve the performance of the octree generation. Morton codes are used as keys to the hash table. [26, 12]. The works by Sundar et al. [16] and Burstedde et al. [19] fall under this category.

There are several ways to construct the Morton code. In our approach, Morton codes are generated recursively from the octree hierarchy and carry all the information about the octree with them for each node of the octree. We explain the details in section 3.

In this part of our implementation, we make use of hash maps, which is a well-known associative container data structure in the field of computer science.

A hash map uses a hash function to produce an index for each key. An efficient implementation of a hash map `unordered_map` is readily available as part of the C++ standard library. C++ `unordered_map` class, is an associative container. It associates every key to a mapped value. The mapped values are stored in buckets in such a way that the elements with key values close to each other are also stored in nearby buckets, which is good for performance. When solving a PDE, each key is mapped to a pointer to the solution vector, which is dependent on the numerical formulation and the discretization method. For example, for a cell-centered finite-volume scheme, this pointer will point to an array of unknowns in each cell, whereas in the case of a finite-element method it will point to the unknowns at the nodes of each element.

The hash map data structure is schematically illustrated in Fig. 5. (We use the term hash table to refer to the fundamental data structure, and hash map to refer to the C++ class.) The only restriction of a hash map is that the key has to be unique. The uniqueness condition is readily satisfied for the present application since every node has a unique Morton code. The only disadvantage of the hash map associative container is that rehashing might occur if the final size of the container is not well predicted. We note that rehashing is an expensive operation computationally. If the size of the array in hash map is big enough data-locality can be preserved and perfect hashing would be possible. However, this is hard to achieve in practice for dynamic problems.

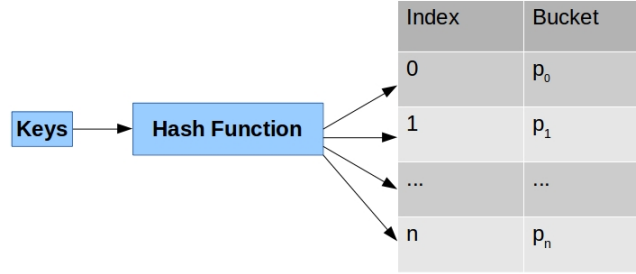


Figure 5: Illustration of hash map data structure, each key is mapped to an index by hash function and each index corresponds to a pointer held in a bucket

Theoretically, a hash table can have $O(1)$ time for insertion, removal and search operations over its elements. However, this is only possible when there are no collisions, which is known as perfect hashing. Let us explore the conditions for perfect hashing in the realm of linear octree representation. First of all, the hash function maps each key to an integer index. It is a common practice to use the remainder operator in the hash function to fit a given number of elements to an array, because a hash table uses array as its underlying data structure. Note that the operation here is limited by the maximum integer number representable on the computer architecture at hand. Consider the mesh given in Fig. 6(a). In order to have a perfect hash, we first convert the Morton code to integer to produce an index and then allocate enough memory. For this simple case the

size of the array should be at least 13

To illustrate the collision issue with the hash table approach, we assume that the initial size of the hash table is 8. Now let us try to fit the 7 elements shown in Fig. 6(a) in the 8 slots by using the remainder operator in the hash function. For elements with hash values of 4 ($4 \bmod 8 = 4$) and 12 ($12 \bmod 8 = 4$) the calculated index is the same. This implies a collision will occur in the fifth slot. In other words, both elements are assigned to the same slot in the hash table. To conform to the Z-order curve and have zero collisions at the same time, the size of the array should be at least 13, which is a substantial increase in allocated memory. We also need to allocate extra space in the array for possible refinements that might occur in the next step. Obviously, this remedy to avoid collision in the hash table can be prohibitively expensive in terms of memory when a mesh with a deep-level of adaptation is desired. Subsequently, one has to relax the no-collision condition and try to use the remainder operator to fit the mesh elements in the array. Once the no-collision condition is relaxed, $O(1)$ access time is no longer guaranteed. Another draw back of using a hash table is that rehashing might occur if one does not predict the final octree size with a good approximation, and rehashing is computationally expensive. Any collision will also break the Z-order curve compliance, leading to performance loss. To have a Z-order curve with a regularly increasing index (i.e., 0, 1, 2, ...) a sorting process is inevitable, which is not readily available in a hashed linear octree generation method. To address this issue we propose to use the red-black tree in place of the hash table.




Figure 6: Illustration of mesh element indexing that arise in the hash table data structure, (a) Morton code representation of a mesh with two levels of adaptation, (b) Hash values corresponding to the Morton code from (a)

2.3. Binarized octree generation

Self-balancing binary search trees, such as a red-black tree, can be favored over a hash table in certain applications, because the worst-case performance

of a red-black tree for insertion, removal and search operations is guaranteed to be $O(\log N)$. A red-black tree also keeps its elements sorted as opposed to the out-of-order storage of elements in a hash table. Red-black trees are widely used in different areas of the computer science field [37]. Its applications ranges from computational geometry to process scheduling in Linux kernels. It guarantees that the path from the root to the farthest node is no more than twice as long as the path from the root to the nearest node [35, 38]. In the present study, we propose to adopt the red-black tree in place of the hash table in the design of the AMR algorithm. Our goal is to create an AMR algorithm that does not lose its Z-order curve compliance at deep levels of adaptation.

A red-black tree is a self-balancing binary search tree with an extra bit for defining the color of a node. The color can be either red or black for a node. The color bit is needed to keep the tree balanced. A red-black tree obeys the following rules:

- The root node is always black
- A red node can not have a red child
- Every path from a given node to any of its descendant nodes contains the same number of black nodes

The memory usage of a red-black tree is much less than an explicit octree. For a large mesh, data-locality due to strict compliance with the Z-order curve offsets the $O(\log N)$ element access cost of a red-black tree, as we have demonstrated in [33].

An associated container that uses a red-black tree data-structure is readily available in the C++ standard template library (STL) as `Set/Map` classes. `Set` class is a special case of the `Map` class in which the associated value is the key itself. The red-black tree needs a comparison function to insert or remove elements. We use the Morton code representation of the octree nodes as the comparison criteria for insertion of the new elements. Note that to compare two keys represented by bits, conversion of the key to an integer number is not required anymore with a red-black tree. We make use of this feature of the red-black tree in conjunction with the `bitset` class in C++ STL to generate meshes with deep levels of adaptation. We use the fast bitwise operation (xor) to perform this comparison.

By adopting a red-black in the AMR algorithm, we accomplish the following:

1. Remove the limitation on maximum level of adaptation. With a red-black tree there is no need to convert the Morton code to an integer index. Red-black tree only requires the comparison of the bits.
2. Preserve the order of the Z-curve exactly and ensure data-locality in the computer implementation for improved performance.

2.4. Graph-based unstructured mesh approach

The Cartesian adaptive mesh generation can also be implemented using a graph-based data structure by treating the mesh as an unstructured grid topology. In this method, the connectivity for each element has to be stored. For

a large mesh with a deep level of adaptation this approach requires a lot of memory because of the need to store connectivity. We have presented an unstructured approach in sufficient detail in [36] and released it as an open-source software [39]. Here, we will make use of it to provide a performance perspective to the other two approaches that we presented above.

3. AMR infrastructure methods

When numerically solving a PDE on an adaptively refined mesh, the AMR library needs to provide certain mesh information, such as connectivity to neighboring elements, coordinates of the elements' centroid and vertices. In this section, we explain the methods to compute such information nimbly by manipulating the Morton code representation of an octree. The methods presented in this section are common to both the hashed linear and binarized octree generation methods.

3.1. Adaptation level

We need to know the current adaptation level of the element under consideration as well as the neighboring elements when constructing a 2:1 balanced octree. The Morton code representation of the octree helps a great deal to extract the level of adaptation without extra storage, because we only need to check the non-zero triple digit.

As an example, let us again consider the node 23 in Fig. 4. The Morton code representation of node 23 is given in Fig. 3. On a 16-bit system, Morton code would be as shown in Fig. 7. Note that the Morton code is padded to the right with zeros to create a 16-bit system representation. Starting from the rightmost digit and scanning to the left, the last non-zero element occurs at level 3, therefore, the level of adaptation for this element is 3. Note that in the present study, element level 0 is the root element and we do not include it in the bit representation.

001,000,101,000,000,0

Figure 7: Morton code for node 23 in Fig. 3 on a prototypical 16-bit system.

For the special case of an element represented with all zeros as Morton code, we need to employ a trick. This special case of all zero bit element becomes important while enforcing the 2:1 balance. Because all digits are zeros, we cannot extract the adaptation level, but we can manipulate the siblings to extract this information. Note that the maximum `bitset` instance size is preset before the refinement as this is a template class where the size of the object is the template parameter. One can start by flipping any one of the digits belonging to the maximum size. Note that by flipping the bit we are now operating on the sibling now. After this we check for the existence of such a key. If it exists then the element under consideration and the sibling belong to the same level.

3.2. Edge length

Edge length is required to calculate volume and surface integrals. Once the adaptation level is known following the above procedure, it is straightforward to calculate the edge length resulting from an isotropic refinement in three-dimensional space, which is simply a generalization of one-dimensional bisection procedure.

Given the edge length at the top level l_0 , the length of the segment at each adaptation level i can be calculated as follows:

$$l_i = \frac{l_0}{2^i}, i = 1, 2, \dots, nLevel. \quad (1)$$

3.3. Centroid coordinates

In the solution of fluid flow equations, whether a collocated or staggered arrangement of the variables is pursued, the centroid of the element needs to be known. Given the initial element's centroid and length of the edges of the initial element in x, y and z directions, it is possible to calculate the centroid of the current element as follows:

$$\begin{aligned} x_i &= x_i^0 - ((-1)^{bit[3j+i]} \frac{1}{2^{j+2}}) l_i, \\ j &= 0, 1, \dots, nLevel - 1, \\ i &= 1, 2, 3. \end{aligned} \quad (2)$$

where x_i^0 implies the x_i component of the centroid of the initial element. Note that summation is implied on the j index. To illustrate with an example, the coordinates of the centroid for the Morton code given in Fig. 4 are calculated as follows:

$$\begin{aligned} x_c &= x_0 + l_0 \left(-\frac{1}{4} - \frac{1}{8} + \frac{1}{16} \right), \\ y_c &= y_0 + l_0 \left(-\frac{1}{4} - \frac{1}{8} - \frac{1}{16} \right), \\ z_c &= z_0 + l_0 \left(+\frac{1}{4} - \frac{1}{8} + \frac{1}{16} \right). \end{aligned} \quad (3)$$

3.4. Element vertices

With an element's centroid and edge length at hand, the vertices of each element is obtained by adding and subtracting the half of the edge length to the centroid at each direction as follows:

$$\begin{aligned}
x &= x_c \pm 0.5 \times l_0, \\
y &= y_c \pm 0.5 \times l_0, \\
z &= z_c \pm 0.5 \times l_0.
\end{aligned} \tag{4}$$

3.5. Connectivity to neighbors

Connectivity to surrounding elements is essential information when solving PDEs, whether it is for calculating the flux of a quantity for enforcing conservation principles or constructing the Jacobian (or stiffness) matrices for implicit schemes. The way connectivity information is stored and accessed in computer memory affects the overall performance tremendously. In this section, we discuss the construction of connectivity that is unique to our implementation. We consider a two-dimensional example for ease of illustration, but our implementation is for the three-dimensional case. Before proceeding further, we define our terminology.

3.5.1. Siblings

Siblings of each element, can be found by flipping the corresponding bit in x, y, z directions, respectively. In this operation, only one bit is flipped. Flipping the bit in this context implies the modification of the bit to *true* if it is *false*, and to *false* when it is originally set to *true*. An example is shown in Fig. 8. For the element highlighted in red, the Morton code in two-dimensional space, using the procedure explained in previous sections, can be constructed as (10,01,01). The siblings for this element can be identified by flipping the corresponding bit in x and y directions. Therefore, the Morton code for the element located in the right hand side is (10,01,11) and for the element underneath is (10,01,00). The siblings are depicted in orange and yellow in Fig. 8, respectively. The element highlighted in cyan is a non-neighbor sibling.




Figure 8: Illustration of neighbor identification from Morton code. Morton codes for different colors are: (00,11,11) for green; (10,01,01) for red; (01,10,00) for yellow; (10,01,11) for orange; (11,00) for magenta; (10,00,10) for cyan.

3.5.2. Off-branch neighbors

The method we presented in section 2.1 for storing an octree as Morton codes enables the identification of off-branch neighbors by only manipulating the bits. Here, the word *off-branch* refers to the neighbors that are not siblings to a given element. Finding non-local siblings in each direction, in this case the green and purple element in Fig. 8, consists of two steps. In the first step, we need to detect the level at which the bit was flipped at the corresponding direction. The technique to accomplish this task is presented in Algorithm 1.

Algorithm 1 Detection of bit-flip level

```

define  $T$  : set of octree Morton codes
given  $M \in T$ 
define  $k = 1, 2, 3$  : for (x, y, or z) directions, respectively
define  $d$  : dimension of the computational space
define  $f$  : flip-level
define  $M_l$  : level of  $M$ 
Initialize :  $\text{boolean} = M(d \cdot l + k)$ 
for  $i = M_l - 1 : 0$  do
    if  $\text{boolean} \neq M(d \cdot i + k)$  then
         $f = i$ 
        break;
    end if
end for
Return  $f$ 

```

Let us describe it with an example. Consider the element highlighted in red in Fig. 8. Morton code for this element can easily be calculated as (10,01,01). Our goal is to construct the Morton code depicted in green without an explicit search. This requires displacement in the x -direction (i.e., $k = 1$ in 1). There are two steps involved in this task:

- The location of the red element in Fig. 8 at the finest level is 01, so the bit in x -direction is 0 (or false). Moving up one level at the tree towards parents, the location of this element is found as 01, which corresponds to false as well. Proceeding one level further, the bit corresponding to that location becomes 1 (true). So this level is the bit-flip level. Therefore, moving from left to right, the bits corresponding to the x -direction is (1,0,0). So the level at which flip has occurred goes all the way up to the first level.
- The second step is to use the information embedded in the bit representation such that the code for the off-branch neighbor is directly extracted. To obtain the Morton code for off-branch neighbor (i.e., the green element in Fig. 8) in x -direction, we flip all the bits in the x -direction up to the bit-flip level. This operation yields the Morton code as (00,10,11).

this element is one of the neighbors. To construct the second neighbor all we have to do is find the sibling of this code in x -direction, which is (10,01,11).

In three-dimensional space, two more neighbors are added to the list. These neighbors are constructed by alternating the bit in x - and z - directions.

3.6. Detecting boundary elements

Boundary elements are those elements that only include siblings as part of their neighborhood connectivity. Boundary elements do not possess any off-branch neighbors. We make use of bitwise operations to identify boundary elements for a given Morton code. The following theorem can be used to detect boundary elements from their Morton code representation.

Theorem 3.1. *An element is a boundary element at a given direction if and only if its Morton code is free of bit-flips in that direction.*

Elements with Morton codes (10,11), (10,10) do not have any off-branch neighbors in the x -direction. Hence, they are boundary elements. If we traverse in the y -direction, elements with codes (10,00) and (10,10) are also boundary elements. The advantage of our approach is that no bits are used for tagging the boundary elements since this criteria can efficiently be extracted from the Morton code at hand.

4. Balancing the octree

During the process to generate an octree with dynamic adaptation, the 2:1 balance may break. Such imbalance propagates, affecting immediate neighbors and neighbors of neighbors. This issue has been coined as the *ripple effect* in the literature of balanced octree generation [16]. Our computer implementation led to the algorithm described in 2. It starts with an initial given queue, and extends that queue to maintain the balance for the entire tree.

There are several ways to implement the $NB \notin Q$ section of the code. We refer the reader to Algorithm 2 for definitions of NB and Q . The $NB \notin Q$ operation can be prohibitively expensive if not handled in a proper way. In the process of balancing, if any of the neighbors of the element under consideration violates the condition, before adding this element to the end of queue, one needs to check if that element is already in the refinement list. When using array data structure or the `Vector` class, the cost of this operation can be calculated to grow as $O(N * N)$, where N is the number of elements tagged for refinement. The cost is prohibitively large for a mesh on the order of several hundred million elements. To alleviate this problem and reduce the $O(N * N)$ cost to $m * O(N)$, where m is the number of iterations done by in the while loop due to the ripple effect, we use a hash map for this list. We note here that for both the hashed linear octree and the binarized octree generation methods, a hash table data structure is used to store the list of the elements tagged for refinement. We

Algorithm 2 2:1 balanced refinement-list generation

```
define  $P = 1$ 
define  $Q$  : List of elements tagged to be refined
define  $NB$  : Neighbor
define  $NB_l$  : Level of neighbor
define  $C_l$  : Level of C element
define  $istart = 0$ ;  $iend = Q.size()$ 
while  $P > 0$  do
  define  $P = 0$ 
  for  $i = istart : iend$  do
     $C = Q(i)$ 
    if  $NB_l < C_l$  &  $NB \notin Q$  then
      append  $NB$  to  $Q$ 
      set  $P = 1$ 
    end if
  end for
   $istart = iend$ ;  $iend = Q.size()$ 
end while
Return  $Q$ 
```

achieve significant reduction in time spent for the balance algorithm with this modification. Timings for the Stanford bunny geometry using the `Vector` and `Map` in search are presented in Table 1. Once again, this example illustrates the value of selecting the right data structure in scientific computing.

Table 1: Comparison of total runtime for two different data structures and their effect on the search algorithms in the 2:1 balance, the Stanford bunny geometry is used in this analysis

level	9	10	11	12	13
mesh size	118,357	437,557	1,199,290	2,830,815	5,867,800
time [s](Hash Map)	0.8	2.57	4.71	10.47	26.0
time [s](Vector)	1.35	5.40	19.02	63.0	185.37

5. Tagging elements for refinement

The mesh generation and adaptation process requires the identification of the elements which intersects the immersed geometry. The brute-force approach to obtain the list of tagged elements for refinement would be through a search. However, a brute-force approach can become expensive for a large mesh. With a brute-force search, the computational cost grows with $m \times n$, where n is the number of points of the geometry and m number of mesh elements at the current level. and ensure data-locality in the computer implementation

One way to alleviate this bottleneck and reduce the cost of search is to voxelize the geometry and perform the search on the specific voxel instead of the entire geometry. This method improves search proportional to number of voxels. Alternatively, by taking advantage of the Morton code representation, it is possible to completely remove this search stage using *geometry encoding*. Removing the search here implies dropping the cost from $O(m \times n)$ to $O(1)$ if hash table used as an underlying container for the list of tagged elements.

Coordinates of the centroid of the STL triangles are used to determine if an element resides in a given voxel. Normals of the surfaces can also be assigned in a similar fashion.

5.1. Voxelization of immersed geometry

Voxelization of the immersed geometry comes handy to restrict certain numerical operations around the geometry. One could consider voxelization as a refined bounding box wrapping around the geometry. The map values are pointers to the geometry file portion that is segmented by the given element. We store the solid geometry in STereoLithography (STL) file format.

The level of refinement in voxelization is much lower than the solution mesh. It is specified independent of the solution according to the solid geometry at hand. The voxel class inherits from the base tree class as shown in Listing 1. We note that voxel need not be a 2:1 balanced octree.




Figure 10: Voxelized representation of the Stanford bunny geometry.

Voxelized view of the Stanford bunny is presented in Fig. 10. One does not need this many voxels to accelerate the search. A higher refinement level is adopted here for illustration purposes.

Listing 1: Base tree class and inherited voxel class

```
template <size_t N, typename value>
class Tree
{
    template <size_t N1, typename value>
    friend class Hdf5Xmf; //write out in hdf5 format

    protected:
        real    rootlength[3]; // length of the first generation
        real    rootcoords[3]; // centroid of the first generation
        bitset<N> rootkey = 0; // root value is always set as 0
        uint    npx; // discretization in x direction
        uint    npy; // discretization in y direction
        uint    npz; // discretization in z direction
        // select either unordered_map or map
    private:
        unordered_map<bitset<N>, value*, hash_override<N>>mesh;
        map<bitset<N>, value*, compare_override<N>>mesh;
        unordered_map<bitset<N>, int> refinelist;
        unordered_map<bitset<N>, int> derefinelist;

    public:
        Tree(real *length, real *coords, uint nx, uint ny, uint nz);
}

template <size_t N, typename value>
class Voxel : public Tree<N, value>
{
    private:
        uint maxlevel; //maximum level of refinement
        uint numMax; //number of elements
        unordered_map<bitset<N>, value*>mesh;

    public:
        Voxel(real *length, real *coords)
        // more member functions
}
```

The voxelization mesh has a similar structure as the solution mesh except that it need not be balanced. Ancestor of each element can be found using the bitwise operations. Therefore, the box that includes the corresponding geometric data can be identified quickly. In flow computations around complex geometry, face normals of the solid geometry are also needed for the calculation of fluxes in finite-volume/discontinuous Galerkin methods or implementation of reconstruction schemes in the immersed boundary method. Voxelization is very useful in retrieving such geometric information. If one is only interested in mesh generation without any solution, the more efficient method of geometric encoding using bisection search is recommended, which we discuss next.

5.2. Geometry encoding

Each point of the geometry may be encoded in a similar fashion as the mesh elements. Given the coordinates of each point and the maximum level of adaptation, it is possible to generate a Morton code for each point at the highest level of adaptation. In other words, the finest level element that withholds the

point is detected and saved in the class tree. This is basically equivalent to using the well-known bisection search in three-dimensional space and is only performed once regardless of the maximum level of adaptation. This way the search stage for tagging elements is eliminated altogether. Geometry encoding is illustrated for a one-dimension case in Algorithm 3. The three-dimensional case can be encoded by using bit interleaving in each directions. In this algorithm, we accumulate true or false for each time we bisect a line and generate a code for a given point.

Algorithm 3 Geometry encoding in one-dimension

```

define maxLevel : maximum desired level of refinement which is a preset
value
define Xmin : minimum x-coordinate of the computational domain
define Xmax : maximum x-coordinate of the computational domain
define xc : center coordinate
define xp : x-coordinate of a given points
define b : instance of std::bitset
for i = 0 : maxLevel do
    xc =  $0.5 \times (Xmax + Xmin)$ 
    if xp < xc then
        b(i) = false
        Xmax = xc
    else
        b(i) = True
        Xmin = xc
    end if
end for

```

The expense to encode the geometry is $O(N)$, where N is the number of triangular elements used to represent the solid geometry that we store as an STL file. The clear advantages of geometry encoding are that it does not grow as the solution mesh grows in size, and it is also cheaper in terms of time and memory than generating another coarse tree for voxelization. Figure 11 shows the improvement in runtime due to geometry encoding and compare the runtime against the voxelization approach.

6. Results

We present examples of 2:1 balanced octree generation for several STL geometries. The test cases include: Stanford bunny, which is a well-known test case from the computer graphics community; a sports car; a wind turbine; complex terrain of Mount Hood, OR. AMR for these geometries are presented in Fig. 12, demonstrating the versatility of the current meshing capability.




Figure 11: Effect of geometry encoding on overall execution time as a function of number of Cartesian mesh elements. Stanford bunny geometry with 88,754 surface triangles is used for this analysis

Table 2: Total mesh size for different adaptation levels

level	10	12	14	16
Bunny (no. of triangles = 88,754)				
No. of elements [millions]	0.43	2.83	10.52	23.96
Sports car (no. of triangles = 844,804)				
No. of elements [millions]	0.11	1.84	16.18	74.09

6.1. Performance analysis

In this section, runtime and memory consumption of the hashed linear octree and binarized octree generation methods are investigated. We also demonstrate deep-level adaptation for the wind turbine geometry using the binarized octree method. Deep-level implies refinement levels that are greater than 21. All the calculations are carried out on a single compute-node of a Linux cluster with 64 GB of DDR3 1600 MHz memory and an Intel Xeon E5-2670 2.60 GHz processor.

The number of elements at different levels for the Stanford bunny and the sports car are presented in Table 2. These numbers are independent from the data structure utilized in the AMR algorithm.

Here, using the two Stanford bunny and a sports car geometries, the runtime for binarized octree and hashed linear octree are presented for levels 10, 12, 14 and 16 in Tables 3 and 4, respectively. The percentage of slow-down due

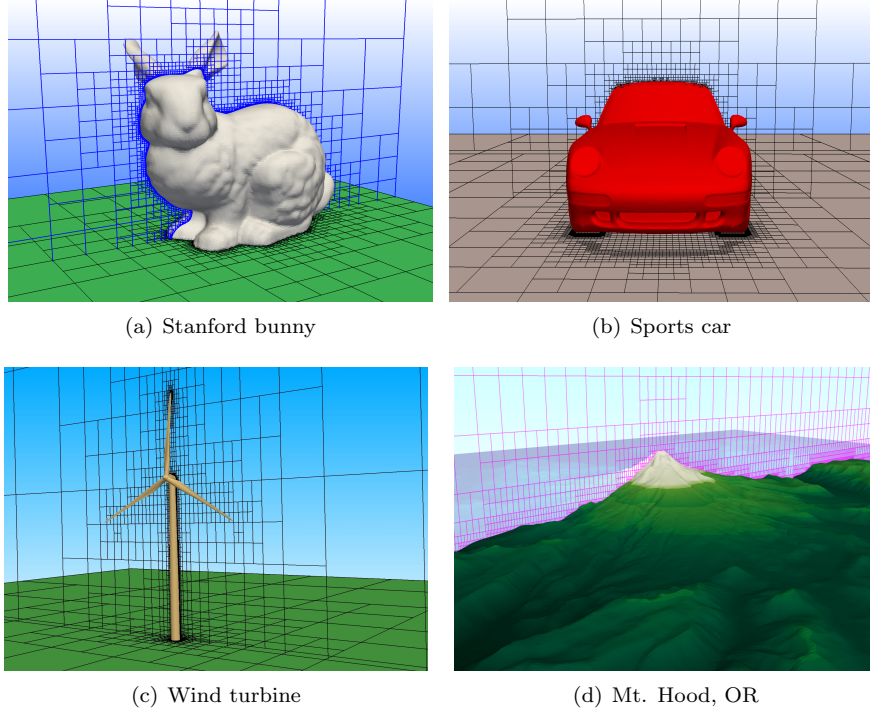


Figure 12: Illustration of the 2:1 balanced Cartesian adaptive mesh generation for complex geometries from different fields.

to the use of a hash table is also presented for further clarification. For both of the geometries, the Z-compliant red-black tree is superior in terms of total runtime.

Table 3: Comparison of total run-time. The Stanford bunny geometry is used for this analysis

level	10	12	14	16
Bunny (No. of triangles = 88,754)				
Z-curve compliant Binarized octree [seconds]	3.3	21.14	69.75	145.87
Z-curve enhanced Hashed linear octree [seconds]	3.7	26.85	87.85	180.01

Memory consumption for the binarized octree and the hashed linear octree are presented for levels 12, 10, 14 and 16 in Tables 5 and 6 for the same geometries discussed in the previous section. Valgrind is used for memory profiling [40]. According to the results, hashed linear octree generation method use less memory than the binarized octree method, but the difference in memory usage

Table 4: Comparison of total run-time. The sports car geometry is used for this analysis

level	10	12	14	16
Car (No. of triangles = 844,804)				
Z-curve compliant binarized octree [seconds]	1.58	19.68	164.56	696.54
Z-curve enhanced hashed linear octree [seconds]	1.64	24.73	260.88	753.58

is not substantial enough to cause any concern.

Table 5: Comparison of memory consumption. The Stanford bunny geometry is used for this analysis.

level	10	12	14	16
Bunny (Ntriangles = 88,754)				
Z-curve compliant binarized octree [MB]	61.4	281.8	792.9	1500
Z-curve enhanced hashed linear octree [MB]	59.6	261.7	675.5	1300
extra memory usage [%]	3.0	7.6	17.4	15.4

Table 6: Comparison of memory consumption. The sports car geometry is used for this analysis.

level	10	12	14	16
Car (No. of triangles = 844,804)				
Z-curve compliant binarized octree [MB]	106.3	341.1	1800	5800
Z-curve enhanced hashed linear octree [MB]	106.3	305.7	2000	5100
extra memory usage [%]	0.0	11.6	-11.1	13.7

Unstructured mesh approaches use much more memory than their tree-based AMR counterparts. The total run-time using puAMR [36] on a single processor is compared against the performance of the binarized octree generation method in Table 7. In this comparison, we are not able to go beyond 16 levels of adaption, because the unstructured approach runs out of computer memory. The other important factor to mention here is that in the unstructured approach, the search time for tagging elements for refinement is much longer than the binarized octree generation that use geometric encoding for the search algorithm. The slow performance of the unstructured approach is expected and we include it in this study to highlight the superior performance of tree-based approaches

for AMR.

Table 7: Comparison of runtime for binarized octree against the graph-based unstructured mesh approach. Stanford bunny geometry is used for this analysis.

level	7	8	9	10	11
mesh size	8163	31242	118357	437,557	1,199,290
Z-curve compliant					
Binarized octree[seconds]	0.16	0.21	0.438	2.57	4.71
Graph-based[seconds]	3.16	4.71	9.02	32.69	615.12

Table 8 demonstrates the deep-level adaption that is inherently feasible with the binarized octree generation approach. We use the wind turbine geometry with 70,406 surface triangles. The mesh around this geometry is presented in Fig. 12(c).

Table 8: Number of mesh elements as a function of adaptation level for a binarized octree generation method. Wind turbine geometry with 70,406 surface triangles is used. Note that hashed linear octree method limits the maximum level of adaption to 21 on 64-bit systems.

level	20	21	22	23
mesh size	1,246,330	2,925,973	6,007,583	10,714,012

7. Conclusions

There are various algorithms for adaptive mesh refinement. Performance of these different approaches can vary greatly. In the present work, we proposed the binarized octree generation method by combining Morton encoding with a red-black tree, which is special kind of self balancing binary tree. In our computer implementation, we benefited from the C++ STL `bitset` `Set/Map` classes. The C++ STL `bitset` enables us to use an arbitrary length for the bit representation regardless of the hardware architecture. Additionally, the red-black tree operates on comparison of bits only. Because of these two features, we are able generate meshes with arbitrary levels of adaptation on a single processor. On the contrary, the hashed linear octree generation method is restricted to 21 levels for 64-bit systems, beyond which an overflow is inevitable because of conversion to an integer number in the hash function. However, one can relax the Z-order curve compliance and generate arbitrary orders of adaptation at the expense of data-locality and loss of runtime performance.

Another advantage of using the red-black tree in octree generation is that it keeps its elements sorted by default and preserves the Z-order curve exactly. In the hashed linear octree generation method, the order can get broken because of collisions in the hash table. An important implication of this strict compliance

with the Z-order curve is that the generated mesh is ready to be partitioned for parallel computations without any need for sorting its elements.

We exploited the bitwise representation of an octree and developed techniques to extract neighborhood connectivity from the Morton codes without explicitly storing the connectivity information. Additionally, we developed a geometry encoding technique to rapidly tag the elements for refinement.

Our results show that the binarized octree technique is faster than the hashed linear octree technique at the expense of a slight increase in memory usage. The current work is released as open-source code which includes both the binarized and the hashed linear octree generation methods [34].

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 1440638 and Army Research Office under Grant No. W911NF-17-1-0564.

References

- [1] I. Babushka, J. Chandra, J. E. Flaherty, Adaptive computational methods for partial differential equations, vol. 16, SIAM, 1983.
- [2] M. J. Berger, J. Oliger, Adaptive mesh refinement for hyperbolic partial differential equations, *J. Comput. Phys.* 53 (3) (1984) 484–512.
- [3] M. Berger, P. Colella, Local adaptive mesh refinement for shock hydrodynamics, *J. Comput. Phys.* 82 (1) (1989) 64 – 84.
- [4] BoxLib, <https://ccse.lbl.gov/BoxLib/>, 2017.
- [5] P. MacNeice, K. M. Olson, C. Mobarry, R. De Fainchtein, C. Packer, PARAMESH: A parallel adaptive mesh refinement community toolkit, *Comput. Phys. Commun.* 126 (3) (2000) 330–354.
- [6] BoxLib, <https://ccse.lbl.gov/AMReX/>, 2017.
- [7] M. Adams, Chombo software package for adaptive mesh refinement applications design document .
- [8] G. L. Bryan, M. L. Norman, B. W. O’Shea, T. Abel, J. H. Wise, M. J. Turk, D. R. Reynolds, D. C. Collins, P. Wang, S. W. Skillman, et al., Enzo: An adaptive mesh refinement code for astrophysics, *Astrophys. J.* 211 (2) (2014) 19.
- [9] SAMRAI, <https://computation.llnl.gov/projects/samrai>, 2017.
- [10] R. D. Hornung, S. R. Kohn, Managing application complexity in the SAMRAI object-oriented framework, *Concurr. Comput.* 14 (5).

- [11] M. J. Berger, R. J. LeVeque, Adaptive Mesh Refinement using Wave-Propagation Algorithms for Hyperbolic Systems, *SIAM J. Numer. Anal.* 35 (1998) 2298–2316.
- [12] T. Lewiner, V. Mello, A. Peixoto, S. Pesco, H. Lopes, Fast generation of pointerless octree duals, in: *Comput. Graph. Forum*, vol. 29, 1661–1669, 2010.
- [13] A. Dubey, A. Almgren, J. Bell, M. Berzins, S. Brandt, G. Bryan, P. Colella, D. Graves, M. Lijewski, F. Lffler, B. OShea, E. Schnetter, B. V. Straalen, K. Weide, A survey of high level frameworks in block-structured adaptive mesh refinement packages, *J. Paral. Distrib. Computing* 74 (12) (2014) 3217 – 3227.
- [14] D. Meagher, Geometric modeling using octree encoding, *Comput. Gr. Image Process.* 19 (2) (1982) 129–147.
- [15] H. Samet, The quadtree and related hierarchical data structures, *ACM Comput. Surv.* 16 (2) (1984) 187–260.
- [16] H. Sundar, R. S. Sampath, G. Biros, Bottom-up construction and 2: 1 balance refinement of linear octrees in parallel, *SIAM J. Sci. Comput.* 30 (5) (2008) 2675–2708.
- [17] I. Gargantini, An Effective Way to Represent Quadtrees, *Commun. ACM* 25 (12) (1982) 905–910.
- [18] I. Gargantini, Linear octrees for fast processing of three-dimensional objects, *Comput. Gr. Image Process.* 20 (4) (1982) 365–374.
- [19] C. Burstedde, L. C. Wilcox, O. Ghattas, p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.* 33 (3) (2011) 1103–1133.
- [20] T. Isaac, C. Burstedde, O. Ghattas, Low-cost parallel algorithms for 2: 1 octree balance, in: *Parallel & Distributed Processing Symposium (IPDPS)*, 2012 IEEE 26th International, IEEE, 426–437, 2012.
- [21] N. Ishikawa, D. Sasaki, K. Nakahashi, Large-scale distributed computation using building-cube method, in: *49th AIAA Aerospace Sciences Meeting*, 754, 2011.
- [22] T. Ishida, S. Takahashi, K. Nakahashi, Efficient and robust cartesian mesh generation for building-cube method, *Journal of Computational Science and Technology* 2 (4) (2008) 435–446.
- [23] T. Tu, D. R. Ohallaron, Balance refinement of massive linear octrees, *Tech. Rep.*, Technical Report CMU-CS-04-129, Carnegie Mellon School of Computer Science, 2004.

- [24] S. F. Frisken, R. N. Perry, Simple and efficient traversal methods for quadtrees and octrees, *J. Graph. Tools* 7 (3) (2002) 1–11.
- [25] M. Schwarz, H.-P. Seidel, Fast parallel surface and solid voxelization on GPUs, in: *ACM Trans. Graph.*, vol. 29, ACM, 179, 2010.
- [26] J. Baert, A. Lagae, P. Dutré, Out-of-core construction of sparse voxel octrees, in: *Comput. Graph. Forum*, vol. 33, Wiley Online Library, 220–227, 2014.
- [27] Tero Karras, Thinking Parallel III, <https://devblogs.nvidia.com/parallelforall/thinking-parallel-part-iii-tree-construction-gpu/>, 2012.
- [28] P. M. Campbell, K. D. Devine, J. E. Flaherty, L. G. Gervasio, J. D. Teresco, Dynamic octree load balancing using space-filling curves, Williams College Department of Computer Science, Tech. Rep. CS-03-01 .
- [29] J. E. Flaherty, R. M. Loy, M. S. Shephard, B. K. Szymanski, J. D. Teresco, L. H. Ziantz, Adaptive local refinement with octree load balancing for the parallel solution of three-dimensional conservation laws, *J. Paral. Distrib. Comput.* 47 (2) (1997) 139–152.
- [30] P. Brunet, I. Navazo, Solid representation and operation using extended octrees, *ACM Trans. Graph.* 9 (2) (1990) 170–197.
- [31] H. T. Vo, C. T. Silva, L. F. Scheidegger, V. Pascucci, Simple and efficient mesh layout with space-filling curves, *J. Graph. Tools* 16 (1) (2012) 25–39.
- [32] G. M. Morton, A computer oriented geodetic data base and a new technique in file sequencing, International Business Machines Company, 1966.
- [33] J. J. Hasbestan, I. Senocak, A short note on the use of the redblack tree in Cartesian adaptive mesh refinement algorithms, *J. Comput. Phys.* 351 (2017) 473 – 477.
- [34] J. J. Hasbestan, I. Senocak, rebl-AMR: Red-black tree based adaptive mesh refinement for immersed-boundary simulations, https://github.com/GEM3D/Hash_vs_RB, version 0.1, DOI = 10.5281/zenodo.836102, 2017.
- [35] A. Drozdek, Data structures and algorithms in C++, Cengage Learning, 2012.
- [36] J. J. Hasbestan, I. Senocak, A parallel adaptive mesh refinement software for complex geometry flow simulations, in: 23rd AIAA Computational Fluid Dynamics Conference, 3301, 2017.
- [37] L. J. Guibas, R. Sedgewick, A dichromatic framework for balanced trees, in: *Foundations of Computer Science, 1978.*, 19th Annual Symposium on, IEEE, 8–21, 1978.

- [38] B. Schneier, Red-black trees, *Dr. Dobbs's Journal* 17 (4) (1992) 42–46.
- [39] J. J. Hasbestan and I. Senocak , puAMR: Parallel adaptive mesh refinement (AMR) with unstructured connectivity, <https://github.com/GEM3D/AMR/puAMR>, version 1.0.0, DOI = 10.5281/zenodo.571216, 2017.
- [40] N. Nethercote, J. Seward, How to shadow every byte of memory used by a program, in: *Proceedings of the 3rd International Conference on Virtual Execution Environments*, ACM, 65–74, 2007.