

# Scalable Algorithms for Parallel Tree-based Adaptive Mesh Refinement with General Element Types

Dissertation

zur  
Erlangung des Doktorgrades (Dr. rer. nat.)  
der  
Mathematisch-Naturwissenschaftlichen Fakultät  
der  
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Johannes Holke

aus

Köln

Bonn, 22. März 2018





Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät  
der Rheinischen Friedrich-Wilhelms-Universität Bonn

1. Gutachter: Prof. Dr. Carsten Burstedde

2. Gutachter: Prof. Dr. Michael Griebel

Tag der Promotion: 06. Juli 2018

Erscheinungsjahr: 2018



# Summary

In this thesis, we develop, discuss and implement algorithms for scalable parallel tree-based adaptive mesh refinement (AMR) using space-filling curves (SFCs). We create an AMR software that works independently of the used element type, such as for example lines, triangles, tetrahedra, quadrilaterals, hexahedra, and prisms. Along with a detailed mathematical discussion, this requires the implementation as a numerical software and its validation, as well as scalability tests on current supercomputers.

For triangular and tetrahedral elements (simplices) with red-refinement (1:4 in 2D, 1:8 in 3D), we develop a new SFC, the tetrahedral Morton space-filling curve (TM-SFC). Its construction is similar to the Morton index for quadrilaterals/hexahedra, as it is also based on bitwise interleaving the coordinates of a certain vertex of the simplex, the anchor node. Additionally, we interleave with a new piece of information, the so called type. The type distinguishes different simplices with the same anchor node. To store the necessary information of a  $d$ -dimensional simplex, we require 10 bytes per triangle and 14 bytes per tetrahedron, which is only one byte more than used in the classical Morton index for quadrilaterals (9 bytes) and hexahedra (13 bytes). For these simplices, we develop element local algorithms such as constructing the parent, children, or face-neighbors of a simplex, and show that most of them are constant-time operations independent of the refinement level.

With SFC based partitioning it is possible that the mesh elements that are partitioned to one process do not form a face-connected domain. The amount of parallel communication among processes with neighboring domains rises with the number of face-connected components. We prove the following upper bounds for the number of face-connected components of segments of the TM-SFC: With a maximum refinement level of  $L$ , the number of face-connected components is bounded by  $2(L - 1)$  in 2D and  $2L + 1$  in 3D. Additionally, we perform a numerical investigation of the distribution of lengths of SFC segments.

Furthermore, we develop a new approach to partition and repartition a coarse (input) mesh among the processes. Compared to previous methods it optimizes for fine mesh load-balance and reduces the parallel communication of coarse mesh data. We discuss the coarse mesh repartitioning algorithm and demonstrate that our method repartitions a coarse mesh of  $371e9$  trees on 917,504 processes (405,000 trees per process) on the Juqueen supercomputer in 1.2 seconds.

We develop an AMR concept that works independently of the element type; achieving this independence by strictly distinguishing between functions that operate on the whole mesh (high-level) and functions that locally operate on a single element or a small set of elements (low-level). We define an application program-

ming interface (API) of low-level functions and develop the high-level functions such that every element-local operation is performed by a low-level function. Thus, by using different implementations of the low-level API for different meshes, or different parts of the same mesh, we are able to use different types of elements.

Many numerical applications, for example finite element and finite volume solvers, require knowledge of a layer of ghost elements. Ghost elements of a process are those elements that lie on a different process but are (face-)neighbors of a process local element. We discuss a new approach to generate and manage these ghost elements that fits into our element-type independent approach. We define and describe the necessary low-level algorithms. Our main idea is the computation of tree-to-tree face-neighbors of an element via the explicit construction of the element's face as a lower dimensional element. In order to optimize the runtime of this method we enhance the algorithm with a top-down search method from Isaac, Burstedde, Wilcox, and Ghattas, and demonstrate how it speeds up the computation by factors of 10 to 20 achieving runtimes comparable to state-of-the art implementations with fixed element types.

With the ghost algorithm we build a straight-forward ripple version of the 2:1 balance algorithm. This is not an optimized version but it serves as a feasibility study for our element-type independent approach.

We implement all algorithms that we develop in this thesis in the new AMR library `t8code`, using the TM-SFC for simplicial and tetrahedral elements. Our modular approach allows us to reuse existing software, which we demonstrate by using the library `p4est` for quadrilateral and hexahedral elements. In a concurrent Bachelor's thesis by David Knapp (INS, Bonn) the necessary low-level algorithms for prisms were developed. With `t8code` we demonstrate that we can create, adapt, (re-)partition, and balance meshes, as well as create and manage a ghost layer. In various tests we show excellent strong and weak scaling behavior of our algorithms on up to 917,504 parallel processes on the Juqueen and Mira supercomputers using up to  $858e9$  mesh elements.

We conclude this thesis by demonstrating how an application can be coupled with the AMR routines. We implement a finite volume based advection solver using `t8code` and show applications with triangular, quadrilateral, tetrahedral, and hexahedral elements, as well as 2D and 3D hybrid meshes, the latter consisting of tetrahedra, hexahedra, and prisms.

Overall, we develop and demonstrate a new simplicial SFC and create a fast and scalable tree-based AMR software that offers a flexibility and generality that was previously not available.

## Acknowledgements

First of all, I want to thank my advisor Prof. Dr. Carsten Burstedde for his constant support and feedback during the development of this thesis, various helpful discussions and lessons in parallel C programming. I would like to thank Prof. Dr.

Michael Griebel for being the second reviewer of this thesis. Thank you to Prof. Dr. Matthias Kreck for being a great mentor. I appreciate the enlightening atmosphere at the INS and thank all my current and previous colleagues. Special thanks to Jose Alberto Fonseca.

Additional thanks to Tobin Isaac for providing the interface to the MPI-3 shared array functionality in the `sc` library.

I appreciate the financial support by the Bonn International Graduate School of Mathematics (BIGS) and travel support by the Bonn Hausdorff Center for Mathematics (HCM).

I gratefully acknowledge the Gauss Centre for Supercomputing e.V. ([www.gauss-centre.eu](http://www.gauss-centre.eu)) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUQUEEN at Jülich Supercomputing Centre (JSC). Additionally, this research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under contract DE-AC02-06CH11357.

Danke an meine Familie für ihre Unterstützung. Vielen Dank an Felix Boes für viele produktive und lustige Nachmittage und für außergewöhnliche Freundschaft.

Mein größter Dank gilt Clelia. Danke für alles, insbesondere für deine bedingungslose Unterstützung, viel Geduld und dass du immer die richtigen Worte zur Motivation findest. *Non ti ringrazierò mai abbastanza per tutto il bene che mi hai fatto.*

# Contents

<b>1. Introduction</b>	<b>11</b>
<b>2. Adaptive Mesh Refinement</b>	<b>16</b>
2.1. Uniform and adaptive mesh refinement . . . . .	16
2.2. Motivation for AMR . . . . .	17
2.3. Unstructured AMR . . . . .	18
2.4. Block-structured AMR . . . . .	19
2.5. Tree-based AMR . . . . .	19
2.5.1. Overview . . . . .	20
2.5.2. Core algorithms . . . . .	22
2.5.3. Independence of the element type . . . . .	23
<b>3. Space-Filling Curves</b>	<b>26</b>
3.1. Refinement spaces and refinements . . . . .	26
3.2. Space-filling curves . . . . .	29
3.3. The Morton space-filling curve . . . . .	30
3.4. Space-filling curves on forests of trees . . . . .	34
3.5. Partitioning with space-filling curves . . . . .	35
<b>4. The Tetrahedral Morton Index</b>	<b>37</b>
4.1. Mesh refinement on simplices . . . . .	37
4.1.1. Bey's refinement rule . . . . .	38
4.1.2. Removal of hanging nodes using red/green refinement . . . . .	42
4.2. The tetrahedral Morton index . . . . .	42
4.2.1. The reference simplex . . . . .	43
4.2.2. The type and Tet-id of a $d$ -simplex . . . . .	43
4.2.3. Encoding of the tetrahedral Morton index . . . . .	44
4.2.4. A different approach to derive the TM-index . . . . .	47
4.2.5. Properties of the TM-index . . . . .	48
4.2.6. The space-filling curve associated to the TM-index . . . . .	51
4.3. Low-level algorithms on simplices . . . . .	53
4.3.1. The coordinates of a $d$ -simplex . . . . .	54
4.3.2. Parent and child . . . . .	55
4.3.3. Neighbor simplices . . . . .	57
4.3.4. The exterior of the root simplex . . . . .	60
4.3.5. A consecutive index for uniform refinements . . . . .	62
4.3.6. Successor and predecessor . . . . .	65

4.4.	High-level AMR algorithms . . . . .	68
4.4.1.	New . . . . .	69
4.4.2.	Adapt . . . . .	70
4.5.	Performance evaluation . . . . .	70
4.6.	Conclusion . . . . .	71
<b>5.</b>	<b>Connected Components of the TM-SFC</b>	<b>74</b>
5.1.	Proof of Theorem 5.1 . . . . .	74
5.1.1.	From uniform to adaptive meshes . . . . .	80
5.1.2.	From one tree to a forest . . . . .	80
5.2.	Enumeration of face-connected segments . . . . .	80
5.3.	Conclusion . . . . .	81
<b>6.</b>	<b>Coarse Mesh Partitioning</b>	<b>83</b>
6.1.	Tree-based AMR . . . . .	84
6.1.1.	The tree shapes . . . . .	85
6.1.2.	Encoding of face-neighbors . . . . .	86
6.1.3.	Orientation between face-neighbors . . . . .	87
6.2.	Partitioning the coarse mesh . . . . .	88
6.2.1.	Valid partitions . . . . .	89
6.2.2.	Encoding a valid partition . . . . .	93
6.2.3.	Ghost trees . . . . .	94
6.2.4.	Computing the communication pattern . . . . .	94
6.2.5.	Face information for ghost trees . . . . .	99
6.3.	Implementation . . . . .	100
6.3.1.	The coarse mesh data structure . . . . .	102
6.3.2.	Updating local indices . . . . .	103
6.3.3.	<code>Partition_cmesh</code> : Algorithm 6.3.1 . . . . .	104
6.4.	Numerical results . . . . .	104
6.4.1.	How to obtain example meshes . . . . .	104
6.4.2.	Disjoint bricks . . . . .	106
6.4.3.	An example with a forest . . . . .	110
6.5.	Conclusion . . . . .	114
<b>7.</b>	<b>Ghost</b>	<b>115</b>
7.1.	Element face-neighbors . . . . .	117
7.1.1.	(i) Identifying the tree face . . . . .	120
7.1.2.	(ii) Constructing the face element . . . . .	120
7.1.3.	(iii) Constructing $F'$ from $F$ . . . . .	123
7.1.4.	(iv) Constructing $E'$ from $F'$ . . . . .	126
7.1.5.	A note on vertex/edge-neighbors . . . . .	127
7.2.	Half-size face-neighbors . . . . .	127
7.3.	Finding owner processes of elements . . . . .	129
7.3.1.	<code>t8_forest_owner</code> . . . . .	130

7.3.2. Owners at a face . . . . .	130
7.4. The ghost algorithms . . . . .	133
7.5. Optimizing the runtime of <code>Ghost</code> . . . . .	134
7.5.1. The recursive top-down search . . . . .	135
7.5.2. The optimized <code>Ghost</code> algorithm . . . . .	137
7.6. Numerical comparison of the ghost versions . . . . .	139
<b>8. 2:1 Balance</b>	<b>142</b>
8.1. Finding leaf descendants of an element . . . . .	143
8.2. The <code>Ripple-balance</code> algorithm . . . . .	145
8.3. Numerical results . . . . .	147
8.3.1. The test case . . . . .	147
8.3.2. Strong scaling . . . . .	148
8.3.3. Weak scaling . . . . .	149
<b>9. A Numerical Application</b>	<b>153</b>
9.1. The advection equation . . . . .	153
9.1.1. Level-set functions . . . . .	153
9.2. Numerically solving the advection equation . . . . .	154
9.2.1. Hanging faces and face-neighbors . . . . .	156
9.2.2. The CFL number . . . . .	156
9.2.3. The refinement criterion . . . . .	157
9.2.4. Error measurement . . . . .	158
9.3. Handling application data . . . . .	158
9.3.1. Interpolation . . . . .	159
9.3.2. Repartitioning of data . . . . .	159
9.3.3. Ghost exchange . . . . .	159
9.4. Tests on a unit cube geometry . . . . .	160
9.4.1. The 2D test case . . . . .	160
9.4.2. The 3D test case . . . . .	160
9.4.3. Convergence tests . . . . .	161
9.4.4. Large scale tests . . . . .	166
9.4.5. Comparison to uniform meshes . . . . .	170
9.4.6. A test with a larger coarse mesh . . . . .	171
<b>10. Conclusion</b>	<b>175</b>
<b>11. Outlook</b>	<b>178</b>
<b>A. The Low-Level API</b>	<b>180</b>



# 1. Introduction

Numerical simulations for the solution of partial differential equations (PDEs) have a wide range of applications in scientific and industrial computing. Almost all of the common methods, including the finite difference method (FD), the finite element method (FEM), the finite volume method (FV), and the discontinuous Galerkin method (DG) use meshes to discretize a given domain on which to solve a PDE, e.g. [20, 90, 135].

Most commonly, meshes consist of quadrilaterals or triangles in 2D, and hexahedra or tetrahedra in 3D. Triangles and tetrahedra are in general better suited to accurately approximate complex domains [127, 128]. However, with quadrilaterals and hexahedra we can use less elements to model the same domain size and many numerical schemes can exploit a tensor product structure and are thus easier to implement on these element types [113, 134, 136, 144, 159]. In some use cases, especially in industrial engineering applications, both advantages are needed, which motivates the use of hybrid meshes with multiple element types. In this case, it is necessary to use prisms or pyramids to transition between hexahedra and tetrahedra [77, 94, 95, 147]. Less common but also in use are hybrid meshes with two element types, either hexahedra and prisms, or tetrahedra and prisms [45, 76].

One of the most successful tools to improve the performance of numerical simulations is the concept of adaptive mesh refinement (AMR), e.g. [6, 49, 133]. For all numerical methods described above, a finer mesh resolution results in a reduction of the computational error and AMR describes the concept of changing this resolution locally; thus, it is possible to maintain fine resolutions in those parts of the mesh where a smaller error is needed while keeping the mesh resolution coarse elsewhere. For this decision process, we may utilize local error estimators [6, 19, 146, 155]. Hence, with AMR we can reduce the number of mesh elements—and thus the memory footprint and runtime of numerical solvers—significantly compared to uniform meshes, which use the same mesh resolution in the whole mesh [33, 79, 102]. However, using AMR, especially in a parallel high-performance computing (HPC) environment, introduces a substantial overhead in mesh management. Particularly demanding tasks include: mesh refining and coarsening, (re-)partitioning in parallel, the creation of (off-process) ghost elements, random access of mesh elements, and more. Therefore, the implementation of complete AMR frameworks is often outsourced into AMR-specialized software libraries, such as `BoxLib` and its successor `AMReX` [4, 18], `Chombo` [42], `Enzo` [22], `libMesh` [78], `p4est` [25], `ParFUM` [87], `Peano` [149], `PUMI` [68], `PYRAMID` [106], `SAMRAI` [151], `Uintah` [15], and others.

A main challenge for AMR on HPC systems is storing and load-balancing the mesh in parallel, in particular when the mesh changes frequently during the com-

putation. Some examples are the simulation of earth mantle convection [33, 84], blood flow [45, 105], two-phase flow and level-set methods [3, 48, 50, 80, 83, 99], and molecular dynamics [5, 115].

One common method for storing the mesh is unstructured meshing, where the element connectivity is modeled as a graph. Partitioning the mesh among parallel processes is then delegated to graph based partitioning methods such as ParMETIS or Scotch [39, 46, 75]. Well-known open source application codes that use unstructured meshes are FEniCS [93], PLUM [107, 108], OpenFOAM [109], or MOAB from the SIGMA toolkit [140]. While these methods allow for approximately optimal mesh partitions (i.e. in terms of surface-to-volume ratio) and provide maximum flexibility of element-to-element connections, their complexity results in slow run-times of the meshing routines and a high demand in memory.

A different approach to efficiently store and partition meshes are space-filling curves (SFCs) [9, 64, 123, 158]. Instead of heuristically solving NP-hard graph partitioning problems, SFCs are used to approximately solve the problem in linear runtime. Starting with a single coarse element and a recursive refinement rule, prescribing how we may replace an element with finer elements covering the same area, we obtain a tree structure of all constructable elements. The tree’s root is the coarse element. A particular adaptive mesh resulting from the coarse element corresponds to the leaves of a subtree. An SFC is a linear order of these leaves. It allows us to efficiently store elements and element data in a linear array according to these SFC indices. Partitioning to  $P$  processes reduces to splitting up the linear order into  $P$  equally sized pieces. This approach was used for quadrilaterals and hexahedra for example in the `Octor` [144] and `Dendro` [124] codes, which demonstrated scaling to tens of thousands of processes.

The SFC approach has then been extended to multiple coarse elements, forming a forest of trees [10, 34, 134]. With this technique it is possible to model any complex geometry as computational domain [10, 134]. The `p4est` code [25] is a particular scalable example of this approach. `p4est` uses quadrilateral and hexahedral elements with the Morton SFC and was shown to scale to up to 3.1 million processes [27, 103, 122].

In the `sam(oa)`<sup>2</sup> framework tree-based AMR with SFCs has been implemented for triangles using the Sierpinski SFC [96] and was shown to scale to over a hundred thousand MPI ranks [97]. Extensions of the 2D Sierpinski curve to 3D tetrahedra exist [9, 123], but we are not aware of a software utilizing them for tree-based AMR.

We begin this thesis with developing a new SFC index for triangular and tetrahedral elements (simplices) with red-refinement [16], the tetrahedral Morton index (TM-index). Its construction is similar to the Morton SFC index for quadrilaterals/hexahedra [101], as it is also based on bitwise interleaving the coordinates of a certain vertex, the anchor node. Opposed to quadrilaterals/hexahedra, a simplex is not uniquely determined by its anchor node and level. We identify a distinguishing property, the type of a simplex. The TM-index is formed by bitwise interleaving

the anchor node coordinates and type of the simplex and all of its ancestors. To store the necessary information of a simplex, we require only the coordinates of the anchor node plus the type of the simplex, resulting in 10 bytes per triangle and 14 bytes per tetrahedron. We demonstrate that in terms of memory and runtime the TM-SFC is comparable to the Morton SFC for quadrilaterals and hexahedra. We prove upper bounds for the number of connected components of any segment of the TM-SFC in 2D and 3D. The simplicial meshes obtained via red-refinement contain so called hanging nodes, faces and edges. These occur at neighboring elements of different refinement levels, where for example a node of the finer element does not coincide with a node of the coarser element, but lies on an edge/face of it. So far, hanging nodes are relatively uncommon for simplicial meshes. However, since they are being successfully used with quadrilateral/hexahedral AMR [1, 34, 132], we are certain that applications are able to incorporate them in their solvers when using the TM-index.

Continuing, we introduce a new approach for partitioning the coarse mesh among the parallel processes. With coarse mesh partitioning we are able to scale the size and complexity of the geometry representation with the number of parallel processes. Typically, such coarse meshes are generated using external tools, such as for example `Gmsh` [57] or `TetGen` [128]. Opposed to previously used methods, we allow multiple owner processes for each tree. We develop the necessary communication pattern with a minimal number of senders, receivers, and messages. This technique enables us to minimize the parallel communication during (re-)partitioning of the coarse mesh while simultaneously maintaining an optimal load-balance of the leaf elements. We also discuss the identification and communication of ghost trees. These are trees that are not local to a process, but are face-neighbor to a process's local tree.

Another contribution of this thesis is the development of a new algorithm to create the ghost elements of the mesh. The main advantage of our method is that it works independently of the used type of elements and thus readily applies, for example, to triangular, tetrahedral, quadrilateral, hexahedral, and hybrid meshes. The core idea is to construct  $d$ -dimensional tree-to-tree face-neighbor elements by explicitly building the  $(d-1)$ -dimensional face and transforming its coordinates before building the neighbor element from it. We optimize our algorithm by utilizing previously published recursive search patterns [70] and obtain runtimes that are comparable to state-of-the-art implementations with fixed element-type [70].

We implement the methods from this thesis in an element-type independent scalable tree-based AMR library, which we call `t8code`<sup>1</sup>. The core mesh handling algorithms that we discuss are `New` for creating a uniform mesh on a given coarse mesh, `Adapt` for refining and/or coarsening elements, `Partition` for redistributing the elements to maintain a balanced load, `Ghost` for creating a layer of off-process ghost elements, and `2:1 Balance` to modify a mesh such that every element's refinement level differs from the level of any face-neighbor by at most  $\pm 1$ ; see e.g. [34, 70]. We reformulate these algorithms in such a way that they do not explicitly use geomet-

---

<sup>1</sup><https://github.com/holke/t8code>

ric or topological information of the individual elements, thus being independent of the element-type. A key technique is to strictly distinguish between functions that operate on the whole mesh (high-level) and element-local functions (low-level).

We define a collection of low-level functions and develop the high-level functions such that every element-local operation is performed by a low-level function. Thus, by using different implementations of the low-level algorithms for different meshes, or different parts of the same mesh, we are able to use different types of elements or SFCs. We also directly obtain a dimension independent formulation. For example, the difference between 2D computations with triangles and 3D computations with tetrahedra is merely a change of the low-level implementation. The modular approach allows us to reuse existing software for the low-level implementations, which we demonstrate by using the library `p4est` for quadrilateral and hexahedral elements.

We perform extensive strong and weak scaling tests of our methods and verify optimal scalability on up to 917,504 parallel processes on the Juqueen and Mira supercomputers using up to 858e9 mesh elements.

This thesis is organized as follows:

In Chapter 2 we give an overview of currently existing adaptive mesh refinement (AMR) techniques with emphasis on tree-based AMR. We continue this overview in Chapter 3, where we discuss SFCs and their connection with tree-based AMR. We present a novel approach to the theory of discrete SFCs that is suited for the purpose of AMR.

In Chapter 4 we introduce a new SFC, the tetrahedral Morton (TM-) SFC for triangular and tetrahedral red-refinement. The TM-SFC is similar to the Morton SFC for hypercubes in that it takes advantage of bitwise interleaving techniques. This chapter was first published in [29] in 2016 by the Society for Industrial and Applied Mathematics (SIAM).

We continue the discussion of the TM-SFC in Chapter 5, where we prove bounds for the number of connected components of segments of the TM-SFC. We show that the number of connected components in a refinement with elements of maximum level  $L$  is at most  $2(L - 1)$  in 2D and at most  $2L + 1$  in 3D and conclude with a numerical investigation of the distribution of the different numbers of connected components occurring.

At the time of this writing the results from Chapter 5 have been submitted for publication and are available as a preprint [32]. This publication also contains results that were obtained by Burstedde and Isaac. However, Chapter 5 contains those excerpts that are the author's own work.

Chapter 6 is devoted to the discussion of the coarse mesh of trees with particular emphasis of a partition technique for it. We describe an approach for the distribution of the trees to the processes that minimizes communication and maintains an optimal load-balance of forest elements. This chapter was first published in [31] in 2017 by the Society for Industrial and Applied Mathematics (SIAM).

In Chapter 7 we discuss a version of the **Ghost** algorithm to create a layer of ghost elements that works with our framework for arbitrary element types. We then discuss a first non-optimized version of **Balance** to establish a 2:1 balance condition among the elements in Chapter 8. We conclude these two chapters with numerical experiments of **Ghost** and **Balance** on hexahedral and tetrahedral meshes. We demonstrate scaling to the full system of JUQUEEN using 458k MPI ranks. We are currently preparing Chapters 7 and 8 for future publication.

We conclude this thesis by presenting a numerical solver in Chapter 9 in order to demonstrate the usability of the presented AMR algorithms for applications. We solve the advection equation with a finite volume scheme on 2D and 3D non-hybrid and hybrid meshes and show excellent strong scaling on up to 458k MPI ranks.

## 2. Adaptive Mesh Refinement

In this chapter we introduce uniform and adaptive mesh refinement (AMR). We discuss advantages and disadvantages of both methods and motivate the usage of AMR. We carry on by giving a brief overview of unstructured and block-structured AMR, and conclude this chapter with a more detailed elaboration of tree-based AMR.

### 2.1. Uniform and adaptive mesh refinement

Most numerical solvers use meshes to discretize an analytical domain on which to solve a PDE. Common methods are finite differences, the finite element method (FEM), the finite volume method (FV), and the discontinuous Galerkin method (DG). See [20, 90, 135] for an overview.

For all these methods a finer mesh resolution results in a smaller computational error. Thus, in order to increase the accuracy of the numerical solution, the computational mesh is replaced with a finer mesh.

One method is to entirely remesh the whole domain with smaller elements with no particular parent-child hierarchy between the old and new mesh. Another refinement method is to replace mesh elements by smaller elements covering the same area as the initial element. In this case, we have a refinement hierarchy between the meshes. The level of an element denotes the number of refinements that were performed from a certain root mesh to that element.

We distinguish between uniform mesh refinement and adaptive mesh refinement (AMR). The former means that each element is refined to the same level, resulting in an evenly structured mesh of same-level elements. With AMR we refine or coarsen individual mesh elements to different levels according to a given refinement criterion. In this way it is possible to concentrate the mesh elements in areas where the computational error is large and maintain a coarse resolution in those regions where the error is already small. In order to decide where to refine or coarsen the mesh, local error estimators are usually taken into account [6, 19, 146, 155]. See also Figure 2.1 for an example of uniform and adaptive meshes.

When we consider time-dependent solvers, such as for example for the Navier-Stokes equations [40, 83], uniform meshes do not change between time steps in general. With AMR, on the contrary, the mesh often changes from one time step to the next. In this case we also speak of dynamic AMR. Dynamic AMR also occurs in static solvers when, for example, the problem is first solved on a coarse mesh in order to estimate the local error according to which the mesh is then refined and the solver started again.




Figure 2.1.: A uniform mesh (left) and an adaptive mesh (right). In a uniform mesh, all elements have the same size or refinement level, while in adaptive meshes the element's refinement levels can differ.

First applications of AMR came up in the mid-80s [14], and using AMR in parallel started in the late 90s, e.g. [59].

## 2.2. Motivation for AMR

The main advantage of uniform meshes is that they are easy to implement. The uniform structure makes iterating over the mesh and finding neighbor elements, etc., near trivial tasks. In parallel computations, load-balancing is straightforward and does not change in time-dependent simulations.

A disadvantage of uniform meshes is that resolving to higher accuracy needs exponentially more elements. Such high element counts increase the memory demand and runtime of an application. It may also happen that the number of elements needed for a desired accuracy results in meshes that are too large to fit into memory.

AMR overcomes these disadvantages of uniform meshes by design. With adaptive meshes it is often possible to use less mesh elements and thus less memory to reach the same accuracy as with a uniform mesh [56]. This permits simulating problems with AMR to an accuracy that cannot be reached with uniform meshes [33, 79].

Another gain of AMR is that computations are faster in comparison to uniform meshes, or reach an increased accuracy while maintaining the same runtime [102].

The main disadvantage of AMR is an increase in algorithmic complexity. For example, for parallel uniform meshes assignment of elements to the processes can be computed straightforwardly and usually does not change during a computation. Using AMR, the local count of elements on each process changes with each adaptation step, which can be as frequent as every time step in a time-dependent simulation. These changes make it necessary to repartition the mesh in order to load-balance the parallel computation. Likewise, computing ghost elements and iterating through the mesh are straightforward operations for uniform meshes, but pose a significant




Figure 2.2.: Three different AMR methods for refining a quadrilateral mesh along a circular arc. From left to right: unstructured AMR, block-structured AMR, and tree-based AMR. For block-structured AMR we outline the boundaries of the rectangular uniform patches in black. For block-structured and tree-based AMR, color refers to the refinement level of an element. In these two methods hanging nodes regularly occur when two elements of different refinement levels are neighbors.

challenge to implement for adaptive meshes.

Furthermore, in variations of AMR so called hanging nodes, edges, and faces may occur. These are nodes (or edges, or faces) of elements that do not align with all nodes of the neighboring elements. Hanging nodes rather lie on an edge or face of at least one neighbor element as in the middle and right pictures in Figure 2.2. We call meshes with hanging nodes non-conforming. When using AMR with hanging nodes, the numerical solvers have to be adapted to handle them properly. See for example [2, 53, 82, 85, 120].

By design AMR introduces a computational overhead, since an application cannot spend 100% of the computing time solving the actual problem, but has to invest resources in managing the mesh. However, many experiments show that with modern computers the overhead introduced by AMR is significantly smaller than the gain in runtime [27, 50, 121].

In conclusion, using AMR certainly involves more elaborate algorithms, and a significant amount of extra research has to be put into the meshing routines compared to uniform meshes. This is a main reason why AMR routines are often outsourced to external libraries and provides a core motivation to implement the AMR library `t8code` in the context of this thesis.

### 2.3. Unstructured AMR

The three most commonly used methods for AMR are unstructured, block-structured, and tree-based AMR, which we briefly compare now.

In unstructured AMR there is no regular pattern of connections between the elements, meaning that the connectivity relations among elements are arbitrary;



see also Figure 2.2 on the left. These meshes are usually conforming; that is, there are no hanging nodes/faces/edges. Unstructured meshes can adapt arbitrarily well to domain geometries. Therefore, they are often used when the domain is somewhat complex [119].

However, the memory requirements for unstructured meshes are higher than for other types of meshes. This is because for each element a list of its neighbors has to be stored, as well as the coordinates of each vertex in the mesh. Memory usage can be orders of magnitude higher compared to tree-based or block structured AMR, where we store coordinates and connectivities only for the elements of a coarse root mesh of trees that has significantly less elements [7].

Solvers using unstructured meshes often apply graph-based partitioning methods to load-balance the mesh. Among the most common partitioner libraries are ParMETIS [75] and Scotch [39]. Although these methods have been applied to billions of elements on millions of processes, their runtime can be orders of magnitude slower compared to partitioning structured meshes [37, 119, 131].

The most commonly used types of elements are triangles in 2D and tetrahedra in 3D, using the properties of Delaunay triangulations [58, 66, 127], but there also exist unstructured mesh codes that use quadrilateral or hexahedral elements [78, 114].

Commonly used libraries for unstructured AMR are `libMesh` [78], `ParFUM` [87], `PUMI` [68], and `PYRAMID` [106]. Some notable application codes are `FEniCS` [93], `PLUM` [107, 108], `OpenFOAM` [109], and `MOAB` from the `SIGMA` toolkit [140].

## 2.4. Block-structured AMR

In block-structured (or patched) AMR—see [13, 120, 130] for example—the mesh is refined in rectangular patches as in the middle picture in Figure 2.2 resulting in a hierarchy of uniform rectangular meshes. Refinement to level  $l + 1$  is only allowed inside a patch of level  $l$  elements. Thus, with block-structured AMR one can use some of the advantages of a uniform mesh within the patches. This comes at the cost of flexibility of the refinement and possibly using more mesh elements than mathematically required.

Block-structured AMR is mainly used with quadrilateral and hexahedral meshes and some common software libraries and applications using block-structure AMR are `BoxLib` and its successor `AMReX` [4, 18], `Chombo` [42], `Enzo` [22], `SAMRAI` [151], and `Uintah` [15].

## 2.5. Tree-based AMR

In this section we provide an overview over tree-based AMR. It is the AMR technique that we use for the algorithms in this thesis.

We start with describing the main idea of tree-based AMR. We then present the core mesh management algorithms and demonstrate in which order they may be




Figure 2.3.: A refined quadrilateral and the associated refinement tree. Starting with the coarse quadrilateral (the root of the tree), we subdivide it into four children (the second row in the tree). We then further subdivide two of these level 1 children into four level 2 children each (the third row in the tree). The final mesh elements correspond to the leaves of the refinement tree.

called by an application. At the end of this section, we introduce our notion of high- and low-level algorithm and the API for general element types in `t8code`.

### 2.5.1. Overview

The core idea of tree-based AMR is the following: Consider the unit square as computational domain, meshed with a single quadrilateral coarse element. This element represents the root of a refinement tree. We can refine it by replacing it with four children and iterate this operation arbitrarily. Thus, the resulting mesh can be represented as a refinement tree; see Figure 2.3. Assigning each element a unique index, we encode the elements in the tree along a so-called space-filling curve (SFC). We can thus store the elements linearly in an array in order of their SFC index. Replacing an element with its children, or four children with its parent, changes the order only locally. We discuss the theory of space-filling curves in more detail in Chapter 3.

Suppose now that the computational domain is more complicated, for example the wing of an airplane, or the mantle of the earth [27,119]. There are two approaches to model such domains using tree-based AMR.

In the first approach, we completely embed the domain inside a single refinement tree and then refine the elements along the domain boundary up to a desired accuracy. All elements that lie outside of the domain are then not considered in the computation [62,91]. See also Figure 2.4 on the left.

The second approach is to model the domain with an unstructured mesh of coarse elements and then understand each coarse element as the root of one refinement tree, giving rise to a forest of elements [10,133,134].

We call the unstructured mesh of tree roots the coarse mesh, which is created a priori to map the topology and geometry of the domain with sufficient fidelity.

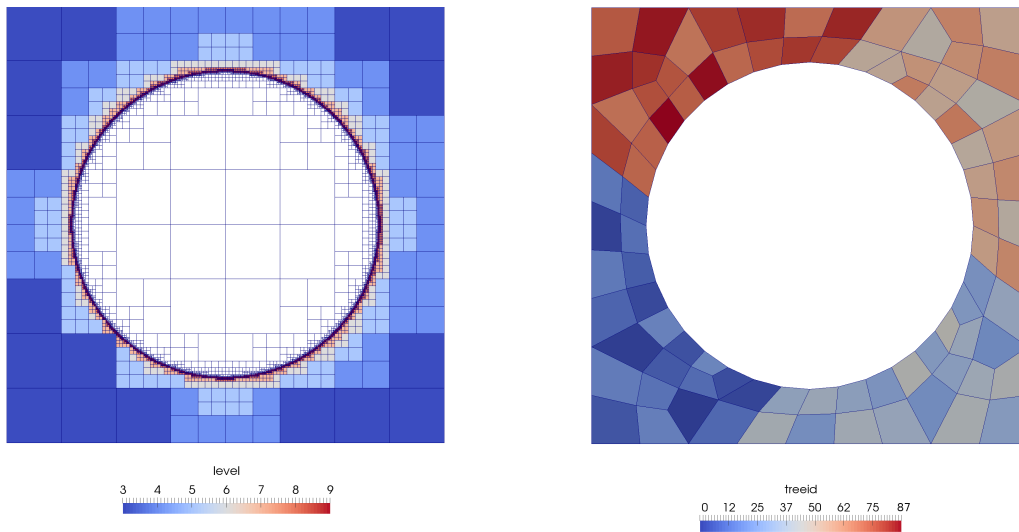


Figure 2.4.: Two ways to approximate a non-trivial domain using tree-based AMR, in this case the region between a circular disk and a unit square. Left: The domain is embedded in a single refinement tree. This tree is then refined at the domain boundary. For a computation only the elements within the domain and intersecting the domain boundary are considered. Color represents the refinement level. Right: The domain is modeled with an unstructured coarse mesh of trees. Each element in this coarse mesh represents one refinement tree. Color represents a consecutive enumeration of the trees. We generated this unstructured mesh with the `Gmsh` mesh generator [57].

Elements may then be refined and coarsened recursively, changing the mesh below the root of each tree. See Figure 2.4 on the right. In extreme cases, for example in industrial and medical applications, the coarse meshes may consist of billions of trees [52, 68, 119]. The accuracy of the approximation of the domain may be further increased by using curved tree edges with the help of higher order geometry functions; see for example [67, 150, 154].

When an application uses the second approach with a large coarse mesh, this mesh has to be partitioned between the processes in order to decrease the memory used by each process. We will discuss coarse mesh partitioning in Chapter 6.

Thus, we have in fact two meshes: The (unstructured) coarse mesh that stores the domain topology and the tree-to-tree connections and the fine mesh that stores the actual computational mesh of refined trees. Usually, the coarse mesh is obtained as the output of a mesh generator, or it is constructed manually for small tree numbers.

Since each coarse mesh element represents a refinement tree, we call the fine mesh forest mesh. Consequently, the tree-based AMR approach is also often called forest-of-trees approach [14, 34]. The forest mesh stores an array of (process local) trees,

and for each tree it stores the linear array of the process local fine elements. We only store the finest elements of a forest and do not store any intermediate elements between the coarse element and the finest elements. Since these elements form the leaves of the refinement trees, we call them leaf elements.

Since each tree stores its own array of elements, the particular linear order (SFC) used is local to the trees. It is thus possible to combine trees of different kinds to form hybrid meshes, for example quadrilaterals and triangles in 2D, or hexahedra, tetrahedra, prisms (and pyramids) in 3D. In theory it is also possible to use different types of linear orders for the same kind of trees, for example Hilbert SFC [65] in one quadrilateral tree and Morton SFC [101] in another. However, we do not discuss such applications in this thesis.

### 2.5.2. Core algorithms

We identify several core algorithms for tree-based AMR.

- **New** – Generate a partitioned uniform mesh on a given geometry.
- **Adapt** – Refine and coarsen mesh elements according to a given criterion.
- **Partition** – Redistribute mesh elements among the processes in order to maintain a balanced load. This has a version for the coarse mesh and one for the fine mesh.
- **Ghost** – Construct and communicate a layer of ghost (halo) elements for each process.
- **Balance** – Establish a 2:1 balance in a mesh. That is, after **Balance**, each mesh element only has neighbors of the same refinement level or at most one level higher or lower.
- **Iterate** – Iterate through the mesh, executing a callback on each element and inter-element interface (faces, edges, or vertices).

We discuss the typical pipeline of these algorithms with a solver application as in Figure 2.5. In a preprocessing step, the geometry of the domain is meshed with a coarse mesh. For small coarse meshes we can construct them in the AMR library itself by explicitly giving the tree-to-tree connections. This approach is not practical for larger coarse meshes, however, which is why we usually use an external mesh generator for this task and feed its output into the AMR library. Among the open source mesh generators we advise the reader to consider **enGrid** [51], **Gmsh** [57], **NETGEN** [125], **TetGen** [128], and **Triangle** [127]. In this step, we may already distribute large coarse meshes among several processes.

After preprocessing, the first step is to create a partitioned uniform forest mesh on top of the coarse mesh with **New**. The initial uniform refinement level depends on the application. Since this mesh is uniform, the global number of elements and the number of elements per tree is known. From this, each process  $p$  can

calculate the index  $E_p$  of its first element and then locally create the elements with indices  $E_p, \dots, E_{p+1} - 1$ . Thus, **New** is completely distributed and does not involve communication.

The next step is usually an initial call to **Adapt** to create a first adapted mesh according to an application's initial refinement criterion. If a 2:1 balance condition is necessary, we call **Balance**. We then often repartition the mesh with **Partition**. This is necessary to maintain a balanced load, since the local element count of several (or all) processes might change during **Adapt** and **Balance**. As a final AMR step, we may call **Ghost** to create a layer of ghost elements.

At this point, the application carries out one or several solver steps, possibly using **Iterate** to iterate through the mesh, for example to identify degrees of freedom to assemble or apply matrices. After solving, we may continue again with **Adapt** using an error estimator and start a new cycle.

Throughout this thesis, we describe theoretical concepts and implementation details of these algorithms. We describe our versions of **New** and **Adapt** in Sections 4.4.1 and 4.4.2. **Partition** is well-understood and we describe it in Section 3.5. **Ghost** and **Balance** are more complex than the other algorithms and therefore we devote a whole chapter for each of them. We describe different versions of **Ghost** in Chapter 7 and discuss **Balance** in Chapter 8. For **Iterate** see our remarks in Chapter 9.

**Remark 2.1.** Typically, the application stores data for each mesh element, for example point values of an interpolated function. If the mesh changes because of adaptation, the data has to be projected onto the new mesh, which for example involves interpolating the function values to new finer mesh elements, or averaging function values of elements that are coarsened into the new value for the parent. See Chapter 9 for a more elaborate discussion.

**Remark 2.2.** The enumeration of the trees in the coarse mesh determines the order in which they are stored and thus their possible partitioning to the processes. The initial enumeration is given by the mesh generator and may be far from optimal for partitioning. For example, if the order is completely random, the partitions on the processes may contain arbitrarily many connected components. In order to reduce the number of connected components and obtain more optimal partitions, we may carry out one run of an unstructured mesh partitioner on the coarse mesh, for example ParMETIS [74] or Scotch [39], as an additional preprocessing step. We discuss the (re-)partitioning of the coarse mesh in detail in Chapter 6.

### 2.5.3. Independence of the element type

We distinguish between global (high-level) algorithms that operate on the complete mesh, such as the core algorithms above, and local (low-level) algorithms that operate on a single element. Examples for the latter type are computing the children or parent of an element, computing the SFC index of an element, or computing face-neighbors of an element.

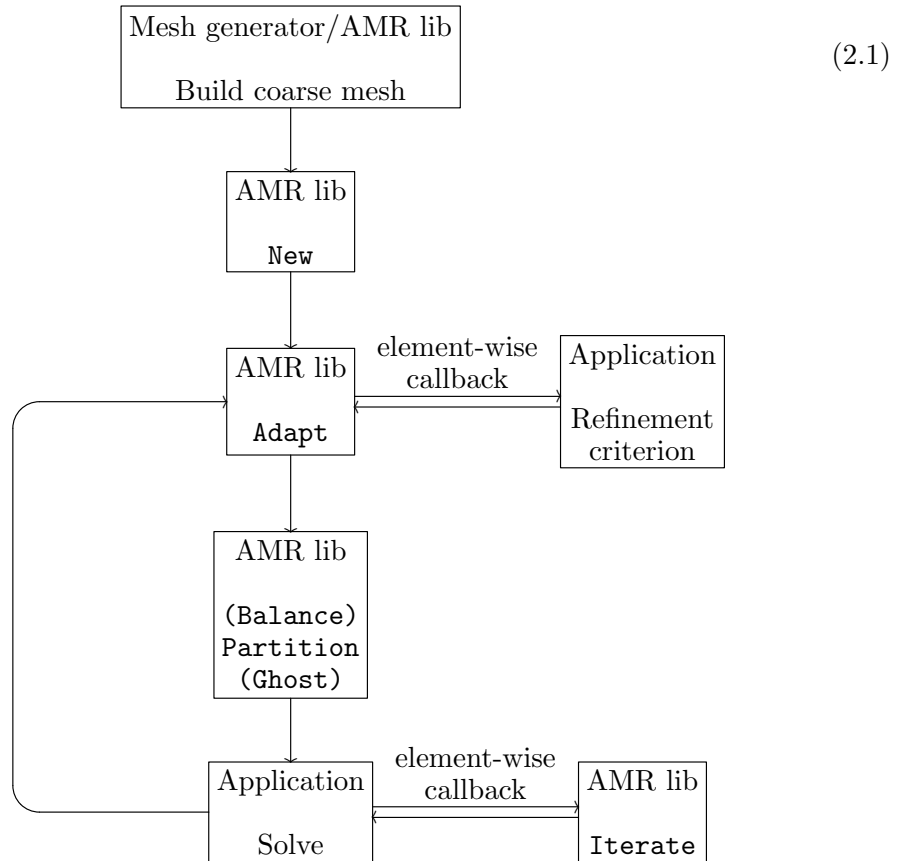


Figure 2.5.: This diagram represents a typical call pipeline for the high-level AMR algorithms of a solver application. At first a coarse mesh is constructed which models the domain geometry. On top of this coarse mesh we construct a distributed uniform forest mesh, which is then adapted according to a refinement criterion provided by the application. We can now balance and partition the mesh and create a layer of ghost elements. Before the application starts the solver. If the solver requires to readapt the mesh, the adapt/balance/partition/ghost cycle starts again.

A main strategy for the algorithms and the code developed in this thesis is the separation of the high-level algorithms from the low-level ones. Take for example the algorithm `Adapt`. In `Adapt`, we iterate over all leaf elements and call a user-provided function for each element and each family of elements. If the function returns positive, we refine the (first) element, thus constructing its children. If the function returns negative on a family of elements, we coarsen them, hence we construct the parent for one of the elements. For an implementation of `Adapt`, however, it is irrelevant how exactly the children or the parent are constructed. We only need to know when to call the corresponding low-level algorithm and can use it as a blackbox. In particular, the implementation of `Adapt` does not change for different types of elements, for example 2D triangles and 3D hexahedra, as long as we provide the appropriate low-level algorithms for the element type.

We carry out the separation of high- and low-level algorithms in the implementation by defining a low-level API, a set of low-level algorithm that define an element type. In `t8code` this low-level API is encoded in form of an abstract C++ base class. Each element type, such as quadrilaterals with Morton SFC index, or tetrahedra with the TM-index, is an implementation of this base class. Hence, introducing a new type of elements of SFC is achieved by simply implementing their low-level algorithms. See for example the Bachelor's thesis [81] in which prisms with a Morton type SFC are implemented. For a detailed list of all low-level functions in `t8code` see Appendix A.

Because of our separation of high- and low-level algorithms, the methods and core algorithms can easily be applied to various kinds of elements and SFCs. Furthermore, mixing elements of different types in the same mesh is possible, allowing us to use hybrid meshes.

## 3. Space-Filling Curves

In this chapter we introduce our notations for space-filling curves and demonstrate their application to tree-based AMR. We discuss how SFCs offer us a way to efficiently store and access the elements of a mesh in an array using the element's SFC index. Furthermore, we show how to partition a mesh in linear runtime to maintain an ideal load-balance using SFCs.

When talking about SFCs, we have to distinguish between SFCs in the analytical sense and SFCs in the numerical (or discrete) sense. In the analytical sense an SFC is a continuous mapping  $f: I \rightarrow \mathbb{R}^n$  from a compact set  $I \subset \mathbb{R}$  into  $\mathbb{R}^n$  whose image  $f(I)$  has a positive  $n$ -dimensional volume [9, 158]. In most cases,  $I$  is the unit interval.

However, when SFCs are used in a numerical sense they are considered as maps from or into a finite index set, which is a discrete version of the analytical SFC. A typical approach is to define an analytical SFC as a limit of some iteration rule and then define an SFC in the numerical sense as a map resulting from only finitely many iteration steps; see for example [9, 64, 123] and the references therein.

In contrast to this indirect approach, we present a new formal definition of (numerical) SFCs that is self-consistent and readily applies to adaptive mesh refinement.

### 3.1. Refinement spaces and refinements

In order to define SFCs, we introduce refinement spaces and refinements. As we have seen in the previous chapter, in tree-based AMR we associate the elements resulting from a root element via refinement with the leaves of a refinement tree. In order to be independent of a particular geometric embedding of the mesh elements, we define SFCs to be mappings on this refinement tree. More precisely, we start with the set of all possible mesh elements, which naturally form the vertices of a tree.

**Definition 3.1.** A **refinement space** is a rooted tree with countably many vertices. Thus, it is a connected graph without circles and a distinguished vertex  $\mathcal{E}$ , the root. We call the vertices of the tree the **elements** of the refinement space.

**Definition 3.2.** The **level**  $\ell(E)$  of an element  $E$  of a refinement space is its distance from the root vertex. By  $\mathcal{S}^l$  we denote the set of all elements of level  $l$ .

The **parent** of an element  $E$  is the unique element  $P$  with which  $E$  shares an edge and for which  $\ell(E) = \ell(P) + 1$ . The root element is the only element that has no parent.

Vice-versa, we say that  $E$  is a **child** of  $P$  if and only if  $P$  is the parent of  $E$ .



**Definition 3.3.** Let  $l \geq 0$ . The  $l$ -th **refinement map**  $R^l$  is a map  $R^l: \mathcal{S}^l \rightarrow \mathcal{P}(\mathcal{S}^{l+1})$ , mapping an element  $E$  of level  $l$  to the set of its children  $R^l(E)$  (here  $\mathcal{P}$  denotes the power set). We say that  $E$  is refined into the elements  $R^l(E)$ .

In fact, we can identify a refinement space solely from its elements, their levels and the refinement maps.

**Proposition 3.4.** A refinement space is equivalent to a triple  $(\mathcal{S}, \ell, \mathcal{R})$ , where  $\mathcal{S}$  is a set,  $\ell: \mathcal{S} \rightarrow \mathbb{N}_0$  a map, and  $\mathcal{R} = \{R^l \mid l \in \mathbb{N}_0\}$  is a set of maps,  $R^l: \mathcal{S}^l \rightarrow \mathcal{P}(\mathcal{S}^{l+1})$  with  $\mathcal{S}^l = \ell^{-1}(l)$ , such that

- there exists exactly one element  $\mathcal{E} \in \mathcal{S}$  with  $\ell(\mathcal{E}) = 0$ , and
- the image of  $R^l$  is a partition of  $\mathcal{S}^{l+1}$ :

$$R^l(E) \cap R^l(E') = \emptyset \text{ for } E \neq E' \in \mathcal{S}^l, \quad (3.1)$$

$$\bigcup_{E \in \mathcal{S}^l} R^l(E) = \mathcal{S}^{l+1}. \quad (3.2)$$

*Proof.* Given such a triple, we build a graph by connecting two elements of  $\mathcal{S}$  if one is the refinement of the other. Properties (3.1) and (3.2) then directly imply that the graph has no circles and is connected. It is thus a tree and  $\mathcal{E}$  is a possible choice for the root. The other direction of the equivalence follows from Definitions 3.2 and 3.3.  $\square$

By abuse of notation we also write  $\mathcal{S}$  instead of  $(\mathcal{S}, \ell, \mathcal{R})$ .

**Remark 3.5.** If a maximum level  $\mathcal{L} = \max_{E \in \mathcal{S}} \ell(E) < \infty$  exists, then we obtain

$$\mathcal{S}^l = \emptyset \text{ for all } l > \mathcal{L}, \quad (3.3)$$

and because of (3.2), we get

$$\mathcal{S}^l \neq \emptyset \text{ for all } l \leq \mathcal{L} \quad (3.4)$$

as well as  $R^{\mathcal{L}}(E) = \emptyset$  for each level  $\mathcal{L}$  element  $E \in \mathcal{S}$ . In particular, such a maximum level always exists if there are finitely many elements.

If a maximum level does not exist, then the set of elements must be infinite and we can also conclude

$$\mathcal{S}^l \neq \emptyset \text{ for all } l \in \mathbb{N}_0. \quad (3.5)$$

The refinement maps define a hierarchy of the elements within a refinement space. We introduce several notations for the relationships among elements.

**Definition 3.6.** If  $E \in R^l(E')$  in a refinement space, then we say that  $E$  is a **child** of  $E'$  and  $E'$  is the **parent** of  $E$ .

Note that because of (3.2) each element except the root element has a parent and because of (3.1) it is unique. It is generally not true that each element of level  $l < \mathcal{L}$  has a child, since we do not exclude the case  $R^l(E) = \emptyset$ .

**Definition 3.7.** We say that  $E$  is a **descendant** of  $E'$  if  $E$  results from  $E'$  via successive refinement, thus  $E = E'$ , or  $\ell(E) \geq \ell(E')$  and  $E \in R^{\ell(E)-1} \circ \dots \circ R^{\ell(E')}(E')$ , Furthermore,  $E'$  is an **ancestor** of  $E$  if and only if  $E$  is a descendant of  $E'$ .

**Definition 3.8.** A **refinement** of a refinement space  $\mathcal{S}$  is a subset  $\mathcal{S}' \subset \mathcal{S}$  of elements that can be constructed from the root element via successively replacing a parent with its children. Thus, starting with the set  $\mathcal{S}^0 = \{\mathcal{E}\}$ , a refinement is a set that can be constructed from it by applying the following rule recursively:

- If  $\mathcal{S}'$  is a refinement and  $E \in \mathcal{S}'$ , then  $\mathcal{S}' \setminus \{E\} \cup R^{\ell(E)}(E)$  is a refinement.

**Remark 3.9.** Thus, in the language of trees, a refinement  $\mathcal{S}'$  consists of the leaves of a subtree of the refinement space that contains the root element and for each element either none or all of its children.

**Definition 3.10.** We call the tree from Remark 3.9 the **refinement tree** of  $\mathcal{S}'$ , and the elements of  $\mathcal{S}'$  are the **leaves** of the refinement.

With this notion, we think of a refinement as an adaptive mesh, as it could arise from any tree-based adaptive mesh application. We display examples for 1:4 quadrilateral refinement in Figures 3.1 and 3.4.

**Remark 3.11.** Each set  $\mathcal{S}^l$  of level  $l$  elements is a refinement, since we can construct it from  $\mathcal{S}^{l-1}$  by applying the refinement rule to each element. This iteration starts with  $\mathcal{S}^0 = \{\mathcal{E}\}$ . We call the refinement  $\mathcal{S}^l$  the **uniform refinement of level  $l$** .

We illustrate this definition with a geometrical example.

**Example 3.12.** As an example we discuss the 1:4 refinement of quadrilateral elements. We fix a maximum level  $\mathcal{L}$  and consider the scaled unit square  $[0, 2^\mathcal{L}]^2$ . Starting with the root element  $\mathcal{E} = [0, 2^\mathcal{L}]^2$ , all other elements in our refinement space are constructed by dividing an element into four subelements as in Figure 3.1, increasing their level by one. The resulting refinement tree is a quadtree; each vertex of level  $l < \mathcal{L}$  has exactly four children. To be formally correct, we describe the refinement space in terms of the level and the refinement maps. Let  $\mathcal{S}^l$  be the set of all subsquares of  $\mathcal{E}$  of side length  $2^{\mathcal{L}-l}$  with coordinates an integer multiple of  $2^{\mathcal{L}-l}$ . We can identify such an element with its lower left corner  $C = (m2^{\mathcal{L}-l}, n2^{\mathcal{L}-l})$  and its level  $l$ . Here,  $m, n \in \mathbb{N}_0$ ,  $0 \leq n, m < 2^l$ . The set  $\mathcal{S}$  is defined as the union of all  $\mathcal{S}^l$ ,  $0 \leq l \leq \mathcal{L}$ . We describe the refinement maps  $R^l$  in terms of the corner coordinates by

$$\begin{aligned}
 R^l((m2^{\mathcal{L}-l}, n2^{\mathcal{L}-l})) = \{ & (2m2^{\mathcal{L}-l-1}, 2n2^{\mathcal{L}-l-1}), \\
 & (2m2^{\mathcal{L}-l-1}, (2n+1)2^{\mathcal{L}-l-1}), \\
 & ((2m+1)2^{\mathcal{L}-l-1}, 2n2^{\mathcal{L}-l-1}), \\
 & ((2m+1)2^{\mathcal{L}-l-1}, (2n+1)2^{\mathcal{L}-l-1}) \}.
 \end{aligned} \tag{3.6}$$




Figure 3.1.: 1:4 refinement of quadrilateral elements. Left: The refinement rule for an element of level  $l < \mathcal{L}$  with anchor node  $C$ . The  $x$  and  $y$  coordinates of  $C$  are integer multiples of  $2^{\mathcal{L}-l}$  and lie within  $[0, 2^{\mathcal{L}}]^2$ . Right: An example for a refinement in the refinement space.

## 3.2. Space-filling curves

In short, an SFC index is a map from a refinement space into the natural numbers that fulfills certain properties.

**Definition 3.13.** A **space-filling curve index** on a refinement space  $\mathcal{S}$  is a map

$$\mathcal{I}: \mathcal{S} \rightarrow \mathbb{N}_0 \quad (3.7)$$

that fulfills the following properties for any  $E, E', \hat{E} \in \mathcal{S}$ :

- (i) The map  $\mathcal{I} \times \ell: \mathcal{S} \rightarrow \mathbb{N}_0 \times \mathbb{N}_0$  is injective. Thus, the index and the level uniquely determine an element of  $\mathcal{S}$ .
- (ii) If  $E$  is an ancestor of  $E'$  then  $\mathcal{I}(E) \leq \mathcal{I}(E')$ . Hence, refining an element cannot decrease its SFC index.
- (iii) If  $\mathcal{I}(E) < \mathcal{I}(\hat{E})$  and  $\hat{E}$  is not a descendant of  $E$ , then  $\mathcal{I}(E) \leq \mathcal{I}(E') < \mathcal{I}(\hat{E})$  for all descendants  $E'$  of  $E$ . Therefore, refining an element is a 'local' operation in terms of the index.

Restricted to the leaves of a refinement, an SFC index becomes injective:

**Proposition 3.14.** Let  $\mathcal{S}$  be a refinement in a refinement space  $\mathcal{S}$  with SFC index  $\mathcal{I}$ . Then any two leaves in  $\mathcal{S}$  have different SFC indices, thus

$$E \neq E' \in \mathcal{S} \Rightarrow \mathcal{I}(E) \neq \mathcal{I}(E'). \quad (3.8)$$

*Proof.* Let  $E \neq E' \in \mathcal{S}$ . Let  $P$  be the nearest common ancestor of  $E$  and  $E'$ , that is, the element in  $\mathcal{S}$  of greatest level that is both an ancestor of  $E$  and  $E'$ . Since the root element is ancestor of all elements,  $P$  must exist and because each element has a unique parent,  $P$  is also unique. Now  $P \neq E$  and  $P \neq E'$  since otherwise  $E'$  would be a descendant of  $E$  or vice-versa, which is a contradiction

with  $\mathcal{S}$  being a refinement. Furthermore, since  $P$  is the nearest common ancestor of  $E$  and  $E'$ ,  $P$  must have two children  $P_E \neq P_{E'}$  that are ancestors of  $E$  and  $E'$ . Since  $\ell(P_E) = \ell(P_{E'}) = \ell(P) + 1$ , we know from the injectivity of  $\mathcal{I} \times \ell$  that the indices of  $P_E$  and  $P_{E'}$  are not equal. We assume  $\mathcal{I}(P_E) < \mathcal{I}(P_{E'})$  without loss of generality. From property (ii) of Definition 3.13 we conclude that  $\mathcal{I}(P_{E'}) \leq \mathcal{I}(E')$  and with property (iii) follows

$$\mathcal{I}(P_E) \leq \mathcal{I}(E) < \mathcal{I}(P_{E'}), \quad (3.9)$$

and thus  $\mathcal{I}(E) < \mathcal{I}(P_{E'}) \leq \mathcal{I}(E')$ , proving the claim.  $\square$

Proposition 3.14 gives us the theoretical justification for the desired one-dimensional storage scheme of a refinement. Since for any refinement the SFC index is injective on the leaves, we can uniquely arrange the leaves in an array such that the order induced by the SFC index is preserved.

**Corollary 3.15.** *Let  $\mathcal{S}$  be a refinement of  $\mathcal{S}$  with SFC index  $\mathcal{I}$  and let  $N = |\mathcal{S}|$  be the number of leaves. Then there exists a unique bijective map*

$$\mathcal{I}_{\mathcal{S}}: \mathcal{S} \rightarrow \{0, \dots, N - 1\}, \quad (3.10)$$

that is monotonous under  $\mathcal{I}$ , thus

$$\mathcal{I}_{\mathcal{S}}(E) < \mathcal{I}_{\mathcal{S}}(E') \Leftrightarrow \mathcal{I}(E) < \mathcal{I}(E'). \quad (3.11)$$

The statement remains true if  $N = \infty$ , in which case the right-hand side of equation (3.10) is  $\mathbb{N}_0$ .

**Definition 3.16.** We call this  $\mathcal{I}_{\mathcal{S}}$  the **consecutive Index** of  $\mathcal{S}$  with regard to  $\mathcal{I}$ .

### 3.3. The Morton space-filling curve

Most SFCs in use arise from a recursive pattern. The root element is a polytope; most common are lines, triangles, quadrilaterals, tetrahedra, hexahedra, prisms, and pyramids. This root element is then subdivided into children of the same polytope type<sup>1</sup>, which are then provided with a local order, i.e. per parent. This refinement pattern is then applied recursively to the children, with the possibility of rotations and/or reflections. See also Figure 3.2.

For a detailed overview of SFCs we refer the reader to [63, 64] and the references therein. Here, we discuss the Morton SFC for quadrilateral/hexahedral elements as an example.

The Morton curve (or  $Z$ -curve) for quadrilateral (2D) and hexahedral (3D) elements was first described by Lebesgue [88] and its applications to data storage were discussed by G. M. Morton [101]. The refinement space for the Morton index

<sup>1</sup>With the exception of pyramids, which may be subdivided into pyramids and tetrahedra.




Figure 3.2.: The child-ids for the Morton index. We label the four children of a quadrilateral with 0, 1, 2, 3 in  $Z$ -order (left). We recursively apply this scheme to all descendants, appending the child-id on each level (right) to compute the Morton index. This is an alternate method to the bitwise interleaving (3.13). See also equation (3.16) and Figure 3.3.

results from 1:4 refinement in 2D and 1:8 refinement in 3D. In general, the Morton index can be defined in any space dimension  $n$  on the  $n$ -dimensional hypercube with  $1:2^n$  refinement. Let us now consider dimension  $n = 2$ . We already discussed in Example 3.12 that each element  $Q$  of the refinement space  $\mathcal{S}$  is identified with its level and its lower left corner coordinates  $(x, y)$ . We call this corner the anchor node of  $Q$ . These coordinates are integers in  $[0, 2^\mathcal{L}) \cap \mathbb{N}_0$  and as such they possess a binary representation of length  $\mathcal{L}$ :

$$x = \sum_{j=0}^{\mathcal{L}-1} x_j 2^j = (x_{\mathcal{L}-1} x_{\mathcal{L}-2} \dots x_0)_2, \quad y = \sum_{j=0}^{\mathcal{L}-1} y_j 2^j = (y_{\mathcal{L}-1} y_{\mathcal{L}-2} \dots y_0)_2, \quad (3.12)$$


with  $x_j, y_j \in \{0, 1\}$ . The Morton index  $m: \mathcal{S} \rightarrow \mathbb{N}_0$  is defined by mapping these coordinates to their bitwise interleaving,

$$m(Q) := (y_{\mathcal{L}-1} x_{\mathcal{L}-1} y_{\mathcal{L}-2} x_{\mathcal{L}-2} \dots y_0 x_0)_2 \in [0, 2^{2\mathcal{L}}). \quad (3.13)$$

See also Figure 3.3. This scheme extends to higher dimensions  $n > 2$  via bitwise interleaving all  $n$  coordinates of the anchor node. The Morton index is an SFC index in the sense of Definition 3.13 [136].

We also discuss a recursive way to describe the Morton index. We label the four children of a single quadrant with 0, 1, 2, and 3 in  $Z$ -order, the so-called child-id. Hence, the lower left child has id 0, the lower right child 1, the upper left child 2, and the upper right child 3; see Figure 3.2.

Each element  $E$  in the refinement space is constructed from the root element via successive refinement. These refinements are unique and thus we obtain a sequence of child-ids  $(c_1, c_2, \dots, c_{\ell(E)})$  describing this refinement process. It can be read as: Start with the root element, take its child  $c_1$ , from this child take the child  $c_2$ , and so forth. The element itself has child-id  $c_{\ell(E)}$  with regard to its parent.



$$\begin{aligned} (x, y) &= (10, 4) = ((1010)_2, (0100)_2) \\ \Rightarrow m(Q) &= (01100100)_2 = (100)_{10} \end{aligned} \quad (3.14)$$

$$m(Q) = (1210)_4 = (100)_{10} \quad (3.15)$$

Figure 3.3.: Two ways of computing the Morton index  $m(Q)$  of a quadrant  $Q$  in the quadrilateral refinement with maximum level  $\mathcal{L} = 4$ . The quadrant's anchor node has coordinates  $x = 10$  and  $y = 4$ . We can compute the Morton index via bitwise interleaving as in equation 3.14. A second way to compute the  $m(Q)$  is via  $Q$ 's refinement path. We can construct  $Q$  from the root element by taking its first child, then the second child of that first child, and finally taking the first child of this quadrant. This leads to the sequence (121) of child-ids. We append zeroes until we reach the length  $\mathcal{L}$  and interpret it as a quaternary number, leading to  $m(Q) = (1210)_4$ .

We can now extend this sequence up to length  $\mathcal{L}$  by filling up with zeroes and since  $c_i \in \{0, 1, 2, 3\}$ , we can interpret it as a quaternary number:

$$(c_1 c_2 \dots c_{\ell(E)} 0 \dots 0)_4 \in [0, 4^{\mathcal{L}}) = [0, 2^{2\mathcal{L}}). \quad (3.16)$$

It is straightforward to show that this is exactly the Morton index of  $E$ , [34, 142]:

$$m(E) = (c_1 c_2 \dots c_{\ell(E)} 0 \dots 0)_4. \quad (3.17)$$

From this we obtain the correspondence

$$(y_i x_i)_2 = c_{\mathcal{L}-i}, \quad (3.18)$$

meaning that we can read of the child-id of  $E$ 's ancestor at level  $i + 1$  from the  $(\mathcal{L} - i)$ -th bits of the anchor node coordinates. See also Figure 3.3. We refer to Figure 3.4 for an example of the SFC arising from the Morton index.

Due to the bitwise interleaving, the Morton index of an element can be computed in constant time and is efficient to implement. Since the index is given implicitly by the anchor node coordinates, most algorithms that operate with the Morton code do not have to carry out the actual bit interleaving but work with the coordinates instead; see for example [34]. Another advantage of the Morton index is that it is memory efficient, storing only the anchor node coordinates and level of an element.




Figure 3.4.: The SFC curve arising from the Morton code for quadrilateral 1:4 refinement. Left: A refinement of the root element together with the SFC starting in the lower left corner and ending in the upper right corner. Right: The associated refinement tree with the SFC passing the leaves (top). Via the consecutive index, the SFC order induces a linear order of the elements in the refinement, such that we can store these elements in an array (bottom).

In addition, many low-level algorithms such as finding an element’s parent, refining an element, or constructing its face-neighbors, etc. can be computed in constant time, independent of the element’s level. Since the logic of the Morton code is the same for 2D and 3D, one can implement both versions with the same source code, carrying out all operations on the  $z$ -coordinate within an `if`-clause or a preprocessor macro, such as it is implemented in the `p4est` library for example [25]. Due to these advantages, the Morton index is chosen as SFC index by various AMR packages and solvers [2, 34, 137].

The greatest perceived disadvantage of the Morton index are the jumps in the SFC. The Morton index does not possess the same locality properties such as for example the Hilbert curve [65], which may lead to an increased runtime when iterating over the mesh elements and can result in cache misses when locating neighbor elements [35]. However, practical experiments show that this effect is not measurable [26]. Also, the Morton SFC may produce disconnected partitions, though the number of connected components is shown to be at most two [9, 32].

It is notable that the Morton index is not only used for AMR applications. Examples for other use cases include fast matrix multiplication with the Strassen-Algorithm [145], image encryption [38], and databases [118].

In Chapter 4 we develop a SFC for triangles and tetrahedra whose core idea is based on the Morton SFC.

### 3.4. Space-filling curves on forests of trees

We do not want to restrict our meshes to single refinement trees, but include the possibility to patch trees together to form more complex coarse meshes. Thus, from the SFC point of view, we have a collection of refinements with individual SFCs on them, which we consider together as a so-called forest.

**Definition 3.17.** Let  $\{\mathcal{K}_0, \dots, \mathcal{K}_{K-1}\}$  be refinements of respective refinement spaces  $\{\mathcal{S}_0, \dots, \mathcal{S}_{K-1}\}$ . Then the **forest**  $\mathcal{F}$  with trees  $\{\mathcal{K}_k\}_{k < K}$  is the set of all leaves of the individual refinements paired with their tree number  $k$ :

$$\mathcal{F} := \bigcup_{k=0}^{K-1} \{k\} \times \mathcal{K}_k. \quad (3.19)$$

The elements of  $\mathcal{F}$  are the **leaves** of the forest.

**Remark 3.18.** The refinement spaces  $\mathcal{S}_k$  do not have to be different from each other. In most cases, all  $\mathcal{S}_k$  are the same or there are only a few different types of refinement spaces.

**Remark 3.19.** A forest with a single tree is isomorphic to a refinement.

**Definition 3.20.** If for a forest  $\mathcal{F}$  each refinement space  $\mathcal{S}_k$  has an SFC index  $\mathcal{I}_k$ , then we extend these to an index  $\mathcal{I}$  on the leaves of  $\mathcal{F}$  by

$$\mathcal{I}: \mathcal{F} \longrightarrow \{0, \dots, K-1\} \times \mathbb{N}_0 \quad (3.20)$$

$$(k, E) \longmapsto (k, \mathcal{I}_k(E)) \quad (3.21)$$

with the order

$$(k, I) < (k', I') :\Leftrightarrow k < k' \text{ or } (k = k' \text{ and } I < I') \quad (3.22)$$

on  $\{0, \dots, K-1\} \times \mathbb{N}_0$ , which extends the individual SFC orders across the trees. By extension of notation we call  $\mathcal{I}$  an SFC index of  $\mathcal{F}$ .

Similarly, the levels of the individual refinement spaces extend to a level map for the forest by  $\ell(k, E) := \ell(E)$ .

With these definitions we form the analogon to Corollary 3.15 for forests.

**Lemma 3.21.** *Let  $\mathcal{F}$  be a forest with SFC index  $\mathcal{I}$  and finite number  $N$  of leaves, then there exists a unique bijective map*

$$\mathcal{I}_{\mathcal{F}}: \mathcal{F} \longrightarrow \{0, \dots, N-1\} \quad (3.23)$$

that is monotonous under  $\mathcal{I}$ , thus

$$(k, \mathcal{I}_k(E)) < (k', \mathcal{I}_{k'}(E')) \Leftrightarrow \mathcal{I}_{\mathcal{F}}(k, E) < \mathcal{I}_{\mathcal{F}}(k', E'). \quad (3.24)$$

We call this map the **consecutive index** of  $\mathcal{F}$ .






Figure 3.5.: SFCs naturally extend to multiple trees to model complex geometries. Here, we show two trees  $k_0$  and  $k_1$  with an adaptive refinement. To enumerate the forest mesh, we establish an a priori order between the two trees and use an SFC within each tree. On the left-hand side of the figure the refinement trees and their linear storage are shown. When we partition the forest mesh to  $P$  processes (here,  $P = 3$ ), we cut the SFC in  $P$  equally sized parts and assign part  $i$  to process  $i$ .

*Proof.* Let  $O_k$  be the number of leaves in all trees of indices smaller than  $k$ :

$$O_k := \sum_{i=0}^{k-1} |\mathcal{K}_k|, \quad (3.25)$$

and let  $\mathcal{I}_{\mathcal{K}_k}$  be the consecutive index of the  $k$ -th tree. Then

$$\mathcal{I}_{\mathcal{F}}(k, E) := O_k + \mathcal{I}_{\mathcal{K}_k}(E) \quad (3.26)$$

satisfies the desired property.  $\square$

We illustrate this concept in Figure 3.5.

### 3.5. Partitioning with space-filling curves

In order to ensure the scalability of an application, it is necessary that each parallel process is assigned (approximately) the same amount of work. This operation is called load-balancing.

While SFCs offer us a way to effectively store and access the leaves of a forest, they also suggest a straightforward method to load-balance these leaves across multiple processes.

Supposing a load-balanced mesh changes due to adaptation, then the workload may not be balanced anymore. It is thus necessary to migrate elements from those processes with more elements to those with less. This procedure is called repartitioning. In order to fairly repartition the mesh, it is necessary to know which elements have to be mapped to which process.

Let  $\mathcal{I}_{\mathcal{F}}$  be the consecutive index of a forest  $\mathcal{F}$  with  $N$  leaves that are to be distributed among  $P$  processes with ranks  $0, 1, \dots, P - 1$ . The partition should respect the SFC order, thus if  $(k, E)$  is assigned to a process  $p$  and  $\mathcal{I}_{\mathcal{F}}(k, E') > \mathcal{I}_{\mathcal{F}}(k', E)$  then we demand that  $(k', E')$  is assigned to a process  $q \geq p$ .

We use the scheme presented in e.g. [34], assigning to process  $p$  the set of leaves

$$\mathcal{F}(p) := \left\{ (k, E) \in \mathcal{F} \mid \left\lfloor \frac{pN}{P} \right\rfloor \leq \mathcal{I}_{\mathcal{F}}(k, E) < \left\lfloor \frac{(p+1)N}{P} \right\rfloor \right\}. \quad (3.27)$$

This ensures that the count of elements on different processes differs by at most one. See also Figure 3.5 for an illustration of the partitioning.

In some cases the computational load differs between elements. We then may assign a non-negative weight  $w(k, E)$  to each element that is proportional to the computational load and demand from the assignment of elements to processes that the sums of all weights of the processes' elements are approximately equal, rather than the count of elements. This is known as the weighted partitioning problem to which several approximative solutions using SFCs exist [34, 47].

An example where weighted partitioning occurs is the *hp*-adaptive finite element method. Here, different elements have different numbers of degrees of freedom such that elements with smaller polynomial degree  $p$  cause less computational load than those with a larger degree [86].

In contrast to load-balancing using graph-based methods [46, 74], the partitions resulting from SFCs can have more surface area and may be disconnected resulting in an increase in parallel communication. However, the partition quality was shown to be acceptable [27] and partitioning with SFC can be orders of magnitude faster, since it reduces the NP-hard problem to an approximation using linear runtime. For this reason SFC partitioning is a common choice when the mesh is repartitioned frequently, e.g. in adaptive solvers for time-dependent PDEs; see for example [9, 23, 60, 62, 123, 156].

## 4. The Tetrahedral Morton Index

This chapter is based on the paper [29]. We edited it slightly in order to fit into the general notations of this thesis, without changing its mathematical content. Copyright © by SIAM. Unauthorized reproduction of this chapter is prohibited.

In this chapter, we develop a new SFC index for triangular and tetrahedral mesh refinement that can be computed using bitwise interleaving operations similar to the Morton index for cubical meshes; see Section 3.3. We demonstrate that this index has many of the favorable properties known for hexahedra. To store sufficient information for random access, we define a low-memory encoding using 10 bytes per triangle and 14 bytes per tetrahedron.

Our starting point is to divide the simplices in a refined mesh into two (two dimensions, 2D), respectively six (three dimensions, 3D), different types and selecting for each type a specific reordering for Bey’s red-refinement [16]. This type and the coordinates of one vertex serve as a unique identifier, the Tet-id, of the simplex in question. In particular, we do not require storing the type of all parent simplices to the root, as one might naively imagine. We then propose a Morton-like coordinate mapping that can be computed from the Tet-id and gives rise to an SFC. Based on this logic, we develop constant-time algorithms (independent of an element’s refinement level) to compute the Tet-id of any parent, child, face-neighbor, and SFC-successor/predecessor and to decide whether for two given elements one is an ancestor of the other. We conclude with scalability tests of mesh creation and adaptation with over 850e9 tetrahedral mesh elements on up to 131,072 cores of the JUQUEEN supercomputer and 786,432 cores of MIRA.

### 4.1. Mesh refinement on simplices

Our aim is to define and examine a new SFC for triangles and tetrahedra by adding ordering prescriptions to the nonconforming Bey-refinement (also called red-refinement) [16, 17, 153]. We briefly restate the red-refinement in this section and contrast it with the well-known conforming (or red/green) refinement.

We refer to triangles and tetrahedra as  $d$ -simplices, where  $d \in \{2, 3\}$  specifies the dimension. It is sometimes convenient to drop  $d$  from this notation. A  $d$ -simplex  $T \subseteq \mathbb{R}^d$  is uniquely determined by its  $d + 1$  affine-independent corner nodes  $\vec{x}_0, \dots, \vec{x}_d \in \mathbb{R}^d$ . Their order is significant, and therefore we write

$$T = [\vec{x}_0, \vec{x}_1, \vec{x}_2] \quad \text{in 2D,} \quad (4.1a)$$

$$T = [\vec{x}_0, \vec{x}_1, \vec{x}_2, \vec{x}_3] \quad \text{in 3D.} \quad (4.1b)$$




Figure 4.1.: Left: The refinement scheme for triangles in two dimensions. A triangle  $T = [\vec{x}_0, \vec{x}_1, \vec{x}_2] \subset \mathbb{R}^2$  is refined by dividing each face at the midpoints  $\vec{x}_{ij}$ . We obtain four smaller triangles, all similar to  $T$ . Right: The situation in three dimensions. If we divide the edges of the tetrahedron  $T = [\vec{x}_0, \vec{x}_1, \vec{x}_2, \vec{x}_3] \subset \mathbb{R}^3$  in half, we get four smaller tetrahedra (similar to  $T$ ) and an inner octahedron. By dividing the octahedron along any of its three diagonals (shown dashed) we finally end up with partitioning  $T$  into eight smaller tetrahedra, all having the same volume. The refinement rule of Bey is obtained by always choosing the diagonal from  $\vec{x}_{02}$  to  $\vec{x}_{13}$  and numbering the corners of the children according to (4.2).

**Definition 4.1.** We define  $\vec{x}_0$  as the **anchor node** of  $T$ . By  $\vec{x}_{ij}$  we denote the midpoint between  $\vec{x}_i$  and  $\vec{x}_j$ , thus  $\vec{x}_{ij} = \frac{1}{2}(\vec{x}_i + \vec{x}_j)$ .

#### 4.1.1. Bey's refinement rule

Bey's rule is a prescription for subdividing a simplex. It is one instance of the so-called red-refinement, where all faces of a simplex are subdivided simultaneously.

**Definition 4.2.** Given a  $d$ -simplex  $T = [\vec{x}_0, \dots, \vec{x}_d] \subset \mathbb{R}^d$ , the **refinement rule of Bey** consists of cutting off four subsimplices at the corners (as in Figure 4.1). In 3D the remaining octahedron is then divided along the diagonal from  $\vec{x}_{02}$  to  $\vec{x}_{13}$ . Bey numbers the  $2^d$  resulting subsimplices as follows.

$$2D : \quad \begin{aligned} T_0 &:= [\vec{x}_0, \vec{x}_{01}, \vec{x}_{02}], & T_1 &:= [\vec{x}_{01}, \vec{x}_1, \vec{x}_{12}], \\ T_2 &:= [\vec{x}_{02}, \vec{x}_{12}, \vec{x}_2], & T_3 &:= [\vec{x}_{01}, \vec{x}_{02}, \vec{x}_{12}], \end{aligned} \quad (4.2a)$$

$$3D : \quad \begin{aligned} T_0 &:= [\vec{x}_0, \vec{x}_{01}, \vec{x}_{02}, \vec{x}_{03}], & T_4 &:= [\vec{x}_{01}, \vec{x}_{02}, \vec{x}_{03}, \vec{x}_{13}], \\ T_1 &:= [\vec{x}_{01}, \vec{x}_1, \vec{x}_{12}, \vec{x}_{13}], & T_5 &:= [\vec{x}_{01}, \vec{x}_{02}, \vec{x}_{12}, \vec{x}_{13}], \\ T_2 &:= [\vec{x}_{02}, \vec{x}_{12}, \vec{x}_2, \vec{x}_{23}], & T_6 &:= [\vec{x}_{02}, \vec{x}_{03}, \vec{x}_{13}, \vec{x}_{23}], \\ T_3 &:= [\vec{x}_{03}, \vec{x}_{13}, \vec{x}_{23}, \vec{x}_3], & T_7 &:= [\vec{x}_{02}, \vec{x}_{12}, \vec{x}_{13}, \vec{x}_{23}]. \end{aligned} \quad (4.2b)$$

**Remark 4.3.** If we apply the refinement rule from Definition 4.2 recursively to the descendants of a  $d$ -simplex  $S$ , we obtain a refinement space in the sense of Definition 3.1 with root element  $S$ . By a refinement of  $S$  we mean a refinement of this refinement space.

Note, that the refinement rule explicitly allows nonuniform meshes and thus hanging faces and edges.

We recall some basic definitions of relations among mesh elements in the simplicial case.

**Definition 4.4.** The  $T_i$  from (4.2) are called the **children** of  $T$ , and  $T$  is called the **parent** of the  $T_i$ , written  $T = P(T_i)$ . Therefore, we also call the  $T_i$  **siblings** of each other. If a  $d$ -simplex  $T$  belongs to a refinement of another  $d$ -simplex  $S$ , then  $T$  is a **descendant** of  $S$ , and  $S$  is an **ancestor** of  $T$ . The number  $\ell$  of refining steps needed to obtain  $T$  from  $S$  is unique and called the **level** of  $T$  (with respect to  $S$ ); we write  $\ell = \ell(T)$ . Usually  $S$  is clear from the context, and therefore we omit it in the notation. By definition,  $T$  is an ancestor and descendant of itself.

Consider the six tetrahedra  $S_0, \dots, S_5 \subset \mathbb{R}^3$  displayed in Figure 4.2. These tetrahedra form a triangulation of the unit cube. The results and algorithms in this chapter rely on the following property [16].

**Property 4.5.** *Refining the six tetrahedra from the triangulation of the unit cube simultaneously to level  $\ell$  results in the same mesh as first refining the unit cube to level  $\ell$  and then triangulating each smaller cube with the six tetrahedra  $S_0, \dots, S_5$ , scaled by a factor of  $2^{-\ell}$  (see Figure 4.3). The same behavior can be observed in 2D when the unit square is divided into two triangles.*

**Remark 4.6.** A key motivation to use the refinement scheme of Bey is that it produces numerically stable meshes. Thus, no matter how much we refine, the mesh elements do not degenerate. The degeneracy of a mesh element  $T$  with volume  $v$  and side lengths  $\{l_i\}$  may be measured by


$$\eta(T) = \frac{12(3v)^{\frac{2}{3}}}{\sum l_i^2} > 0. \quad (4.4)$$

If  $\eta(T) \ll 1$  the element's volume is small relative to the sum of its side lengths. Imagine a very flat tetrahedron, or a triangle with one angle being close to  $\pi$ .

From Property 4.5 we conclude that for each tetrahedron  $T$  that is an ancestor of one of the  $S_i$  we have


$$\eta(T) = \eta(S_j) \quad (4.5)$$

for some  $j \in \{0, \dots, 5\}$ . Hence,  $\eta(T)$  is bounded and any mesh resulting from Bey's refinement rule is numerically stable. See [92] and the references therein for a more thorough discussion of this subject.



2D	$\vec{x}_0$	$\vec{x}_1$	$\vec{x}_2$
$S_0$	$c_0$	$c_1$	$c_3$
$S_1$	$c_0$	$c_2$	$c_3$

(4.3a)



3D	$\vec{x}_0$	$\vec{x}_1$	$\vec{x}_2$	$\vec{x}_3$
$S_0$	$c_0$	$c_1$	$c_5$	$c_7$
$S_1$	$c_0$	$c_1$	$c_3$	$c_7$
$S_2$	$c_0$	$c_2$	$c_3$	$c_7$
$S_3$	$c_0$	$c_2$	$c_6$	$c_7$
$S_4$	$c_0$	$c_4$	$c_6$	$c_7$
$S_5$	$c_0$	$c_4$	$c_5$	$c_7$

(4.3b)

Figure 4.2.: The basic triangle (2D) and tetrahedra types (3D) obtained by dividing  $[0, 1]^d$  into simplices of varying types, denoted by a subscript. Top left: The unit square can be divided into two triangles sharing the edge from  $(0, 0)^T$  to  $(1, 1)^T$ . We denote these triangles by  $S_0$  and  $S_1$ . The four corners of the square are numbered  $c_0, \dots, c_3$  in  $yx$ -order. Top right: The corner nodes of  $S_0$  and  $S_1$  in terms of the square corners. Bottom left (exploded view): In three dimensions the unit cube can be divided into six tetrahedra, all sharing the edge from the origin to  $(1, 1, 1)^T$ . We denote these tetrahedra by  $S_0, \dots, S_5$ . The eight corners of the cube are numbered  $c_0, \dots, c_7$  in  $zyx$ -order (redrawn and modified with permission [16]). Bottom right: The corner nodes of the six tetrahedra  $S_0, \dots, S_5$  in terms of the cube corners.

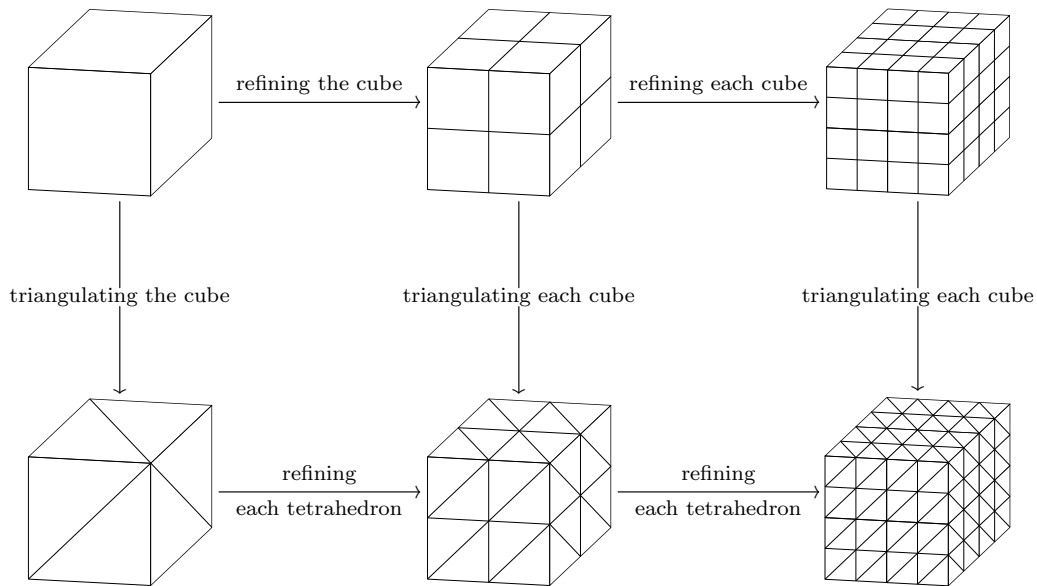


Figure 4.3.: Triangulating a cube according to Figure 4.2 and then refining the tetrahedra via Bey's refinement rule results in the same mesh as first refining the cube into eight subcubes and afterward triangulating each of these cubes. Each occurring tetrahedron is uniquely determined by the subcube it lies in plus its type. The same situation can be observed in 2D if we restrict our view to one side of the cube.




Figure 4.4.: To resolve hanging nodes we can execute an additional step of green- or blue-refinement after the last refinement step [11,36]. Here we show the 2D refinement rules. Left: green (1 hanging node). Right: blue (2 hanging nodes). If a triangle has 3 hanging nodes it is red-refined.

#### 4.1.2. Removal of hanging nodes using red/green refinement

It is worth noting that, although the methods and algorithms presented in this chapter apply to red-refined meshes with hanging nodes, it is possible to augment them to create meshes without hanging nodes. For this we may use red/green or red/green/blue refinement methods [11,36].

After the red-refinement step we may add an additional and possibly nonlocal refinement operation that ensures a maximum level difference of 1 between neighboring simplices. Such an operation is also called 2:1 balance [69,136,144]; we describe it in detail in Chapter 8. Hanging nodes are then resolved by bisecting those simplices with hanging nodes (green/blue refinement) [9, section 12.1.3]. The 2D case is shown in Figure 4.4.

If one of the newly created simplices shall be further refined, the bisection is reversed, the original simplex is red-refined, and the balancing and green-refinement is repeated. This may void the nesting property of certain discrete function spaces, yet applications may still prefer this approach over the manual implementation of hanging node constraints.

## 4.2. The tetrahedral Morton index

As we repeat in Section 3.3, the Morton index or Z-order for a cube in a hexahedral mesh is computed by bitwise interleaving the coordinates of the anchor node of the cube [101]. In this section we present an index for  $d$ -simplices that also uses the bitwise interleaving approach, the tetrahedral Morton index (TM-index). To define the TM-index we look at refinements of a reference simplex, which we discuss in Section 4.2.1 below. For each  $d$ -simplex in a refinement of the reference simplex we define a unique identifier, the so-called Tet-id, which serves as the input to compute the TM-index and for all algorithms related to it. This Tet-id consists of the coordinates of the anchor node of the considered simplex plus one additional number, the type of the simplex. We define the Tet-id and type in Section 4.2.2. We then define the TM-index in Section 4.2.3 and in the following subsections. We show that the TM-index defines a SFC index in the sense of Definition 3.13 and discuss properties of the resulting SFC.



One novel aspect of this construction lies in logically including the types of the simplex and all its parents in the interleaving, while only using the type of the simplex itself in the algorithms.

### 4.2.1. The reference simplex

Throughout the rest of this chapter, let  $\mathcal{L}$  be a given maximal refinement level. Instead of the unit cube  $[0, 1]^d$ , we consider the scaled cube  $[0, 2^\mathcal{L}]^d$ , ensuring that all node coordinates in a refinement up to level  $\mathcal{L}$  are integers. Suppose we are given some  $d$ -simplex  $T \subset \mathbb{R}^d$  together with a refinement  $\mathcal{S}$  of  $T$ . By mapping  $T$  affine-linearly to  $2^\mathcal{L}S_0$  the refinement  $\mathcal{S}$  is mapped to a refinement  $\mathcal{S}'$  of  $2^\mathcal{L}S_0$ . Therefore, to examine SFCs on refinements of  $T$ , it suffices to examine SFCs on  $2^\mathcal{L}S_0$ . Thus, we only consider refinements of the  $d$ -simplex  $T_d^0 := 2^\mathcal{L}S_0$ . Let  $\mathcal{T}_d$  be the set of all possible descendants of this  $d$ -simplex with level smaller than or equal to  $\mathcal{L}$ ; thus

$$\mathcal{T}_d = \{ T \mid T \text{ is a descendant of } T_d^0 \text{ with } 0 \leq \ell(T) \leq \mathcal{L} \}. \quad (4.6)$$

$\mathcal{T}_d$  together with the refinement from Definition 4.2 and  $T_d^0$  as root element is a refinement space in the sense of Definition 3.1.

Any refinement (up to level  $\mathcal{L}$ ) of  $T_d^0$  is a subset of  $\mathcal{T}_d$ , and for each  $T \in \mathcal{T}_d$  there exists at least one refinement  $\mathcal{T}$  of  $T_d^0$  with  $T \in \mathcal{T}$ . In this context, we refer to  $T_d^0$  as the **root simplex**. Furthermore, let  $\mathbb{L}^d$  denote the set of all possible anchor node coordinates of  $d$ -simplices in  $\mathcal{T}_d$ , thus

$$\begin{aligned} \mathbb{L}^2 &= \{ [0, 2^\mathcal{L}]^2 \cap \mathbb{Z}^2 \mid y \leq x \}, \\ \mathbb{L}^3 &= \{ [0, 2^\mathcal{L}]^3 \cap \mathbb{Z}^3 \mid y \leq z \leq x \}. \end{aligned} \quad (4.7)$$

Note that we could have chosen any other of the  $S_i$  (scaled by  $2^\mathcal{L}$ ) as the root simplex and we do not see any advantage or disadvantage in doing so.

### 4.2.2. The type and Tet-id of a $d$ -simplex

Making use of Property 4.5, we define the following.

**Definition 4.7.** Each  $d$ -simplex  $T \in \mathcal{T}_d$  of level  $\ell$  lies in a  $d$ -cube of the hexahedral mesh that is part of a uniform level  $\ell$  refinement of  $[0, 2^\mathcal{L}]^d$ . This specific cube is the **associated cube** of  $T$  and denoted by  $Q_T$ . The  $d$ -simplex  $T$  is a scaled and shifted version of exactly one of the six tetrahedra  $S_i$  that constitute the unit cube, and we define the **type** of  $T$  as this number,  $\text{type}(T) := i$ .

The anchor node of a subcube of level  $\ell$  is the particular corner of that cube with the smallest  $x$ -,  $y$ - (and  $z$ -) coordinates. This means that for each simplex  $T$  in the refinement from Figure 4.3 the anchor node of  $T$  and the anchor node of its associated cube coincide. Any two  $d$ -simplices in  $\mathcal{T}_d$  with the same associated cube are distinguishable by their type.

Ct	Child			
	$T_0$	$T_1$	$T_2$	$T_3$
2D				
b	0	0	0	1
	1	1	1	0

Ct	Child				
	$T_0, \dots, T_3$	$T_4$	$T_5$	$T_6$	$T_7$
3D					
b	0	4	5	2	1
	1	3	2	5	0
	2	0	1	4	3
	3	5	4	1	2
	4	2	3	0	5
	5	1	0	3	4

Table 4.1.: For a  $d$ -simplex  $T$  of type  $b$  the table gives the types  $\text{Ct}(T_0), \dots, \text{Ct}(T_{2^d-1})$  of  $T$ 's children. The corner-children  $T_0, T_1, T_2$  (and in 3D also  $T_3$ ) always have the same type as  $T$ .

From Bey's observation from Figure 4.3 it follows that any simplex in  $\mathcal{T}_d$  can be obtained by specifying a level  $\ell$ , then choosing one level  $\ell$  subcube of the root cube and finally fixing a type. This provides motivation for the following definition.

**Definition 4.8** (Tet-id). For  $T = [\vec{x}_0, \dots, \vec{x}_d] \in \mathcal{T}_d$  we define the **Tet-id** of  $T$  as the tuple of its anchor node and type; thus

$$\text{Tet-id}(T) := (\vec{x}_0, \text{type}(T)). \quad (4.8)$$

**Corollary 4.9.** *Let  $T, T' \in \mathcal{T}_d$ . Then  $T = T'$  if and only if their Tet-ids and levels are the same.*

Note that in an arbitrary adaptive mesh there can be simplices with different levels and each simplex  $T$  has an associated cube of level  $\ell(T)$ . In particular, simplices with the same anchor node can have different associated cubes if their levels are not equal.

Since any simplex in  $\mathcal{T}_d$  can be specified by the Tet-id and level, the Tet-id provides an important tool for our work. The construction of the TM-index in the next section and the algorithms that we present in Section 4.3 rely on the Tet-id as the fundamental data of a simplex. All information about a mesh can be extracted from the Tet-id and level of each element.

Since the root simplex has type 0, in a uniform refinement more simplices have type 0 than any other type. However, a close examination of Table 4.1 together with a short inductive argument leads to the following proposition.

**Proposition 4.10.** *In the limit  $\mathcal{L}, \ell \rightarrow \infty$  the different types of simplices in a uniform level  $\ell$  refinement of  $\mathcal{T}_d$  occur in equal ratios.*

### 4.2.3. Encoding of the tetrahedral Morton index

In addition to the anchor coordinates the TM-index also depends on the types of all ancestors of a simplex. In order to define the TM-index we start by giving a formal definition of the interleaving operation and some additional information.

**Definition 4.11.** We define the **interleaving**  $a \dot{\perp} b$  of two  $n$ -tuples  $a = (a_{n-1}, \dots, a_0)$  and  $b = (b_{n-1}, \dots, b_0)$  as the  $2n$ -tuple obtained by alternating the entries of  $a$  and  $b$ :

$$a \dot{\perp} b := (a_{n-1}, b_{n-1}, \dots, a_0, b_0). \quad (4.9)$$

The interleaving of more than two  $n$ -tuples  $a^1, \dots, a^m$  is defined analogously as the  $mn$ -tuple

$$a^1 \dot{\perp} \dots \dot{\perp} a^m := (a_{n-1}^1, a_{n-1}^2, \dots, a_{n-1}^m, a_{n-2}^1, \dots, a_0^{m-1}, a_0^m). \quad (4.10)$$

**Remark 4.12.** The TM-index of a  $d$ -simplex  $T \in \mathcal{T}_d$  that we are going to define is constructed by interleaving  $d + 1$   $\mathcal{L}$ -tuples, where the first  $d$  are the binary representations of the coordinates of  $T$ 's anchor node and the last is the tuple consisting of the types of the ancestors of  $T$ .

**Definition 4.13.** Let  $T \in \mathcal{T}_3$  be a tetrahedron of refinement level  $\ell$  with anchor node  $\vec{x}_0 = (x, y, z)^T \in \mathbb{L}^3$ . Since  $x, y, z \in \mathbb{N}_0$  with  $0 \leq x, y, z < 2^\mathcal{L}$ , we can express them as binary numbers with  $\mathcal{L}$  digits, writing

$$x = \sum_{j=0}^{\mathcal{L}-1} x_j 2^j, \quad y = \sum_{j=0}^{\mathcal{L}-1} y_j 2^j, \quad z = \sum_{j=0}^{\mathcal{L}-1} z_j 2^j. \quad (4.11)$$

We define the  $\mathcal{L}$ -tuples  $X, Y$ , and  $Z$  as the  $\mathcal{L}$ -tuples consisting of the binary digits of  $x, y$ , and  $z$ ; thus,

$$X = X(T) := (x_{\mathcal{L}-1}, \dots, x_0), \quad (4.12a)$$

$$Y = Y(T) := (y_{\mathcal{L}-1}, \dots, y_0), \quad (4.12b)$$

$$Z = Z(T) := (z_{\mathcal{L}-1}, \dots, z_0). \quad (4.12c)$$

In 2D we get the same definitions with  $X$  and  $Y$ , leaving out the  $z$ -coordinate.

**Definition 4.14.** For a  $T \in \mathcal{T}_d$  of level  $\ell$  and each  $0 \leq j \leq \ell$  let  $T^j$  be the (unique) ancestor of  $T$  of level  $j$ . In particular,  $T^\ell = T$ . We define  $B(T)$  as the  $\mathcal{L}$ -tuple consisting of the types of  $T$ 's ancestors in the first  $\ell$  entries, starting with  $T^1$ . The last  $\mathcal{L} - \ell$  entries of  $B(T)$  are zero:

$$B = B(T) := \left( \underbrace{\text{type}(T^1), \text{type}(T^2), \dots, \text{type}(T)}_{\ell \text{ entries}}, 0, \dots, 0 \right), \quad (4.13)$$

Thus, if we write  $B$  as an  $\mathcal{L}$ -tuple with indexed entries  $b_i$

$$B = B(T) = (b_{\mathcal{L}-1}, \dots, b_0) \in \{0, \dots, d! - 1\}^\mathcal{L}, \quad (4.14)$$

then the  $i$ th entry  $b_i$  is given as

$$b_i = \begin{cases} \text{type}(T^{\mathcal{L}-i}) & \mathcal{L} - 1 \geq i \geq \mathcal{L} - \ell, \\ 0 & \mathcal{L} - \ell > i \geq 0. \end{cases} \quad (4.15)$$

**Definition 4.15** (tetrahedral Morton Index). We define the **tetrahedral Morton index** (**TM-index**)  $m(T)$  of a  $d$ -simplex  $T \in \mathcal{T}_d$  as the interleaving of the  $\mathcal{L}$ -tuples  $Z$  (for tetrahedra),  $Y$ ,  $X$  and  $B$ . Thus,

$$m(T) := Y \dot{\perp} X \dot{\perp} B \quad (4.16a)$$

for triangles and

$$m(T) := Z \dot{\perp} Y \dot{\perp} X \dot{\perp} B \quad (4.16b)$$

for tetrahedra.

This index resembles the well-known Morton index or Z-order for  $d$ -dimensional cubes, which we denote by  $\tilde{m}$  here. For such a cube  $Q$  the Morton index is usually defined as the bitwise interleaving of its coordinates. Thus  $\tilde{m}(Q) = Z \dot{\perp} Y \dot{\perp} X$ , respectively,  $\tilde{m}(Q) = Y \dot{\perp} X$ ; see [34, 101, 136] as well as Section 3.3.

As we show in Section 4.3, the TM-index can be computed from the Tet-id of  $T$  with no further information given. Thus, in an implementation it is not necessary to store the  $\mathcal{L}$ -tuple  $B$ .

The TM-index of a  $d$ -simplex builds up from packs of  $d$  bits  $z_i$  (for tetrahedra),  $y_i$ , and  $x_i$  followed by a type  $b_i \in \{0, \dots, d! - 1\}$ . Since  $d! = 2 < 4$  for  $d = 2$ , we can interpret the 2D TM-index as a quarternary number with digits  $(y_i x_i)_2$  and  $b_i$ :

$$\begin{aligned} m(T) &= ((y_{\mathcal{L}-1} x_{\mathcal{L}-1})_2, b_{\mathcal{L}-1}, \dots, (y_0 x_0)_2, b_0)_4 \\ &= \sum_{i=0}^{\mathcal{L}-1} ((2y_i + x_i)4^{2i+1} + b_i 4^{2i}). \end{aligned} \quad (4.17a)$$

Similarly we can interpret it as an octal number with digits  $(z_i y_i x_i)_2$  and  $b_i$  for  $d = 3$ , since then  $d! = 6 < 8$ :

$$\begin{aligned} m(T) &= ((z_{\mathcal{L}-1} y_{\mathcal{L}-1} x_{\mathcal{L}-1})_2, b_{\mathcal{L}-1}, \dots, (z_0 y_0 x_0)_2, b_0)_8 \\ &= \sum_{i=0}^{\mathcal{L}-1} ((4z_i + 2y_i + x_i)8^{2i+1} + b_i 8^{2i}). \end{aligned} \quad (4.17b)$$

The entries in these numbers are only nonzero up to the level  $\ell$  of  $T$ , since  $x_{\mathcal{L}-i} = y_{\mathcal{L}-i} = (z_{\mathcal{L}-i} =) b_{\mathcal{L}-i} = 0$  for all  $i > \ell$ . The octal/quarternary representation (4.17) directly gives an order on the TM-indices, and therefore it is possible to construct an SFC from it, which we examine further in Section 4.2.6. We use  $m(T)$  to denote both the  $(d+1)\mathcal{L}$ -tuple from (4.16) and the number given by (4.17).

Let us look at Figure 4.5 for a short example to motivate this definition of the TM-index. Since the anchor coordinates and the type together with the level uniquely determine a  $d$ -simplex in  $\mathcal{T}_d$ , one could ask why we do not define the index to be  $((Z \dot{\perp})Y \dot{\perp} X, \text{type}(T))$ , a pair consisting of the Morton index of the associated cube of  $T$  and the type of  $T$ . This index was introduced for triangles in a slightly modified version as semiquadcodes in [112] and would certainly require less information since the computation of the sequence  $B$  would not be necessary. However,




Figure 4.5.: Comparing a straightforward definition of a Morton-type SFC with our approach. Left: The curve arising from taking the Morton order of the quadrants and only dividing into triangles on the last level. Thus the index is  $(Y \dot{\perp} X, \text{type}(T))$ . As we see on the two coarse triangles that are shaded, the children of a level 1 triangle are not necessarily traversed before any other triangle is traversed. Thus, it breaks the locality property that is part of the definition of an SFC, and therefore this index is not suitable for our purposes. Right: The curve arising from the TM-index from our Definition 4.15. We see that for each level 1 triangle its four children are traversed as a group. Theorem 4.18 states that the curve is in fact a proper SFC in the sense of Definition 3.13 and thus the locality property holds for any parent triangle/tetrahedron. The order in which the children are traversed depends (only) on the type of the parent and is different from Bey's order given by (4.2).

it results in an SFC that traverses the leaf cubes of a refinement in the usual Z-order and inside of each cube it traverses the  $d!$  different simplices in the order  $S_0, \dots, S_{d!-1}$ . As a result, there can be simplices  $T$  whose children are not traversed as a group, which means that there is a tetrahedron  $T'$ , which is not an ancestor or descendant of  $T$ , such that some child  $T_i$  of  $T$  is traversed before  $T'$  and  $T'$  is traversed before another child  $T_j$  of  $T$ . Thus, this curve is not an SFC in the sense of Definition 3.13. In contrast to this, Theorem 4.18 states that the TM-index is in fact an SFC-index. Figure 4.5 compares the two approaches for a uniform level 2 refinement of  $T_2^0$ .

#### 4.2.4. A different approach to derive the TM-index

There is another interpretation of the TM-index, which is particularly useful for the AMR algorithms presented in Section 4.3. In order to define it we introduce the concept of the so-called cube-id. According to Figure 4.2 we number the  $2^d$  corners of a  $d$ -dimensional cube by  $c_0, \dots, c_{2^d-1}$  in a  $zyx$ -order ( $x$  varies fastest). When refining a cube to  $2^d$  children, each child has exactly one of the  $c_i$  as a corner, and it is therefore convenient to number the children by  $c_0, \dots, c_{2^d-1}$  as well. For the child




Figure 4.6.: Left: A square is refined to four children, each of which corresponds to a corner of the square. The number of the corner is the cube-id of that child. Right: In three dimensions a cube is refined to eight children. Their cube-ids and corner numbers are shown as well.

$c_i$  we call the number  $i$  the cube-id of that child; see Figure 4.6 for an illustration. Since each cube  $Q$  that is not the root has a unique parent, it also has a unique cube-id. This cube-id can easily be computed by interleaving the last significant bits of the  $z$ - (in 3D),  $y$ -, and  $x$ -coordinates of  $Q$ 's anchor node.

**Definition 4.16.** Because each  $d$ -simplex  $T \in \mathcal{T}_d$  of level  $\ell$  has a unique associated cube we define the **cube-id** of  $T$  to be the cube-id of the associated cube of  $T$ , that is, the  $d$ -cube of level  $\ell$  that has the same anchor node as  $T$ .

If  $X, Y$  (and  $Z$ ) are as is in Definition 4.13 then we can write the cube-id of  $T$ 's ancestors as

$$\begin{aligned} \text{cube-id}(T^i) &= (y_i x_i)_2 && \text{in 2D,} \\ \text{cube-id}(T^i) &= (z_i y_i x_i)_2 && \text{in 3D,} \end{aligned} \quad (4.18)$$

and therefore using (4.17) we can rewrite the TM-index of  $T$  as

$$m(T) = (\text{cube-id}(T^1), \text{type}(T^1), \dots, \text{cube-id}(T^\ell), \text{type}(T^\ell), 0, \dots, 0)_{2^d}. \quad (4.19)$$

This resembles the Morton index of the associated cube  $Q_T$  of  $T$ , since we can write this as

$$\tilde{m}(Q_T) = (\text{cube-id}(Q^1), \dots, \text{cube-id}(Q^\ell), 0, \dots, 0)_{2^d}. \quad (4.20)$$

#### 4.2.5. Properties of the TM-index

In this section we show that the TM-index is an SFC-index. As a first result, we show that  $m \times \ell$  is injective.

**Proposition 4.17.** *Together with a refinement level  $\ell$ , the TM-index  $m(T)$  uniquely determines a  $d$ -simplex in  $\mathcal{T}_d$ .*

*Proof.* If  $\ell = 0$ , then there is only one simplex of level  $\ell$  in  $\mathcal{T}_d$ , which is  $T_d^0$ . So let  $\ell > 0$  and  $m = m(T)$  be given as in (4.16), and let  $\ell$  be the level of  $T$ . From  $m$  we can compute the  $x$ -,  $y$ - (and  $z$ -) coordinates of the associated cube of  $T$ . We can also compute the type of  $T$  from the TM-index. By Corollary 4.9 this information uniquely determines  $T$ .  $\square$

For the Morton index  $\tilde{m}$  for cubes the following important properties are known [136]:

- (i) A Morton index of a cube  $Q$  is the prefix of an index of a cube  $P$  of higher level than  $Q$  if and only if  $P$  is a descendant of  $Q$ .
- (ii) The Morton indices of the descendants of a parent cube are larger than or equal to the index of the parent cube.
- (iii) Refining only changes the SFC locally. Thus, if  $Q$  is a cube and  $P$  is a cube with  $\tilde{m}(Q) < \tilde{m}(P)$  and  $P$  is not a descendant of  $Q$ , then  $\tilde{m}(Q') < \tilde{m}(P)$  for each descendant  $Q'$  of  $Q$ .

Property (iii) defines a hierarchic invariant of the SFC that is specific to our construction (see Figure 4.5). We show below that properties (i), (ii) and (iii) hold for  $d$ -simplices and the TM-index described by (4.16). Proposition 4.17 together with (ii) and (iii) imply that the TM-index is an SFC-index in the sense of Definition 3.13.

**Theorem 4.18.** *For arbitrary  $d$ -simplices  $T \neq S \in \mathcal{T}_d$  the TM-index satisfies the following:*

- (i) *If  $\ell(T) < \ell(S)$ , then  $m(T)$  is a prefix of  $m(S)$  if and only if  $S$  is a descendant of  $T$ .*
- (ii) *If  $T$  is an ancestor of  $S$  then  $m(T) \leq m(S)$ .*
- (iii) *If  $m(T) < m(S)$  and  $S$  is no descendant of  $T$ , then for each descendant  $T'$  of  $T$  we have*

$$m(T) \leq m(T') < m(S). \quad (4.21)$$

The proof of Theorem 4.18 requires some work and we need to show a technical result first. Hereby, we consider only the 3D case, since for 2D the argument is completely analogous. We define an embedding of the set of all TM-indices into the set of Morton indices for 6D cubes. Since the properties (i)–(iii) hold for these cubes it follows that they hold for tetrahedra as well. To this end, for a given tetrahedron  $T \in \mathcal{T}_3$  we interpret each entry  $b_j$  of  $B(T)$  as a 3-digit binary number

$$b_j = (b_j^2 b_j^1 b_j^0)_2, \quad (4.22)$$

which is possible since  $b_j \in \{0, \dots, 5\}$ . We obtain three new  $\mathcal{L}$ -tuples  $B^2, B^1, B^0$  satisfying

$$B = B^2 \dot{\perp} B^1 \dot{\perp} B^0, \quad (4.23)$$

and thus we can rewrite the TM-index as

$$m(T) = Z \dot{\downarrow} Y \dot{\downarrow} X \dot{\downarrow} B^2 \dot{\downarrow} B^1 \dot{\downarrow} B^0. \quad (4.24)$$

Note that we can interpret each  $B^i$  as an  $\mathcal{L}$ -digit binary number for which we have  $0 \leq B^i < 2^\mathcal{L}$ . Now let  $\mathcal{Q}$  denote the set of all 6D cubes that are a child of the cube  $Q_0 := [0, 2^\mathcal{L}]^6$ :

$$\mathcal{Q} = \{Q \mid Q \text{ is a descendant of } Q_0 \text{ of level } 0 \leq \ell \leq \mathcal{L}\}. \quad (4.25)$$

Since a cube  $Q \in \mathcal{Q}$  is uniquely determined by the six coordinates  $(x_0, \dots, x_5)$  of its anchor node plus its level  $\ell$ , we also write  $Q = Q_{(x_0, \dots, x_5), \ell}$ . Note that the Morton index for a cube can be defined as the bitwise interleaving of its anchor node coordinates [101]:

$$\tilde{m}(Q) = X^5 \dot{\downarrow} X^4 \dot{\downarrow} X^3 \dot{\downarrow} X^2 \dot{\downarrow} X^1 \dot{\downarrow} X^0. \quad (4.26)$$

**Proposition 4.19.** *The map*

$$\begin{aligned} \Phi: \mathcal{T}_3 &\longrightarrow \mathcal{Q}, \\ T &\longmapsto Q_{(B^0(T), B^1(T), B^2(T), x(T), y(T), z(T)), \ell(T)} \end{aligned} \quad (4.27)$$

is injective and satisfies

$$\tilde{m}(\Phi(T)) = m(T). \quad (4.28)$$

Furthermore, it fulfills the property that  $T'$  is a child of  $T$  if and only if  $\Phi(T')$  is a child of  $\Phi(T)$ .

*Proof.* The equation  $m(T) = \tilde{m}(\Phi(T))$  follows directly from the definitions of the TM-indices on  $\mathcal{T}_3$  and  $\mathcal{Q}$ . From Lemma 4.17 we conclude that  $\Phi$  is injective. Now let  $T', T \in \mathcal{T}_3$ , where  $T'$  is a child of  $T$ . Furthermore, let  $\ell = \ell(T)$ . We know that  $Q' := \Phi(T')$  is a child of  $Q := \Phi(T)$  if and only if for each  $i \in \{0, \dots, 5\}$  it holds that

$$x_i(Q') \in \left\{ x_i(Q), x_i(Q) + 2^{\mathcal{L} - (\ell + 1)} \right\}. \quad (4.29)$$

Because of the underlying cube structure (compare Figure 4.3) we know that the  $x$ -coordinate of the anchor node of  $T'$  satisfies

$$x(T') \in \left\{ x(T), x(T) + 2^{\mathcal{L} - (\ell + 1)} \right\}, \quad (4.30)$$

and likewise for  $Y(T')$  and  $Z(T')$ . Therefore, (4.29) holds for  $i = 3, 4, 5$ . By definition  $B^j(T')$  is the same as  $B^j(T)$  except at position  $\mathcal{L} - (\ell + 1)$ , where

$$B^j(T')_{\mathcal{L} - (\ell + 1)} = b_{\mathcal{L} - (\ell + 1)}^j(T') \in \{0, 1\} \quad (4.31)$$

and

$$B^j(T)_{\mathcal{L} - (\ell + 1)} = 0. \quad (4.32)$$



Hence, we conclude that (4.29) also holds for  $i = 0, 1, 2$ . So  $\Phi(T')$  is a child of  $\Phi(T)$ .

To show the other implication, let us assume that  $\Phi(T')$  is a child of  $\Phi(T)$ . Since  $\ell(T') = \ell(\Phi(T')) > 0$ ,  $T'$  has a parent and we denote it by  $P$ . In the argument above we have shown that  $\Phi(P)$  is the parent of  $\Phi(T')$  and because each cube has a unique parent the identity  $\Phi(P) = \Phi(T)$  must hold. Therefore, we get  $P = T$  since  $\Phi$  is injective; thus,  $T'$  is the child of  $T$ .  $\square$

Inductively we conclude that  $T'$  is a descendant of  $T$  if and only if  $\Phi(T')$  is a descendant of  $\Phi(T)$ . Now Theorem 4.18 follows, because the desired properties (i)–(iii) hold for the Morton index of cubes [136].  $\square$

#### 4.2.6. The space-filling curve associated to the TM-index

By interpreting the TM-indices as  $2^d$ -ary numbers as in (4.17) we get a total order on the set of all possible TM-indices, and therefore it gives rise to an SFC for any refinement  $\mathcal{S}$  of  $T_d^0$ . In this section we further examine the SFC derived from the TM-index. We give here a recursive description of it, similarly to how it is done for the Sierpinski curve and other cubical SFC by Haverkort and van Walderveen [64].

Part (iii) of Theorem 4.18 tells us that the descendants of a simplex  $T$  are traversed before any other simplices with a higher TM-index than  $T$  are traversed. However, the order that the children of  $T$  have relative to each other can be different to the order of children of another simplex  $T'$ . In particular the order of the simplices defined by the TM-index differs from the order (4.2) defined by Bey. We observe this behavior in 2D in Figure 4.5 on the right-hand side: For the level 1 triangles of type 0 the children are traversed in the order

$$T_0, T_1, T_3, T_2 \tag{4.33}$$

and the children of the level 1 triangle of type 1 are traversed in the order

$$T_0, T_3, T_1, T_2. \tag{4.34}$$

In fact, the order of the children of a simplex  $T$  depends only on the type of  $T$ , as we show in the following Proposition.

**Proposition 4.20.** *If  $T, T' \in \mathcal{T}_d$  are two  $d$ -simplices of given type  $b = \text{type}(T) = \text{type}(T')$ , then there exists a unique permutation  $\sigma \equiv \sigma_b$  of  $\{0, \dots, 2^d - 1\}$  such that*

$$\begin{aligned} m(T_{\sigma(0)}) &< m(T_{\sigma(1)}) < \dots < m(T_{\sigma(2^d-1)}), \\ &\text{and} \\ m(T'_{\sigma(0)}) &< m(T'_{\sigma(1)}) < \dots < m(T'_{\sigma(2^d-1)}). \end{aligned} \tag{4.35}$$

*Thus, the children of  $T$  and the children of  $T'$  are in the same order with respect to their TM-index.*

*Proof.* By ordering the children of  $T$  and  $T'$  with respect to their TM-indices, we obtain  $\sigma$  and  $\sigma'$  with

$$\begin{aligned} m(T_{\sigma(0)}) &< m(T_{\sigma(1)}) < \cdots < m(T_{\sigma(2^d-1)}), \\ m(T'_{\sigma'(0)}) &< m(T'_{\sigma'(1)}) < \cdots < m(T'_{\sigma'(2^d-1)}). \end{aligned} \quad (4.36)$$

These permutations are well-defined and unique with this property because different simplices of the same level never have the same TM-index; see Proposition 4.17. It remains to show that  $\sigma' = \sigma$ . Let  $\ell = \ell(T)$  and  $\ell' = \ell(T')$ . Since the TM-indices of the children of  $T$  do all agree up to level  $\ell$ , we see, using the notation from (4.17), that their order  $\sigma$  depends only on the  $d + 1$  numbers ( $z$  is omitted for  $d = 2$ )

$$z_{\mathcal{L}-(\ell+1)}(T_i), \quad y_{\mathcal{L}-(\ell+1)}(T_i), \quad x_{\mathcal{L}-(\ell+1)}(T_i) \quad \text{and} \quad b_{\mathcal{L}-(\ell+1)}(T_i). \quad (4.37)$$

The same argument applies to  $\sigma'$  and  $\ell'$ . From now on we carry out the computations for  $d = 3$ . Since  $\text{type}(T) = \text{type}(T')$  we can write

$$T = \lambda T' + \vec{c}, \quad (4.38)$$

with

$$\lambda = 2^{\ell'-\ell}, \quad \vec{c} = \begin{pmatrix} x(T) - x(T') \\ y(T) - y(T') \\ z(T) - z(T') \end{pmatrix}. \quad (4.39)$$

Since the refinement rules (4.2) commute with scaling and translation we also obtain

$$T_i = \lambda T'_i + \vec{c} \quad (4.40)$$

for the children of  $T$  and  $T'$  and therefore

$$b_{\mathcal{L}-(\ell+1)}(T_i) = \text{type}(T_i) = \text{type}(T'_i) = b_{\mathcal{L}-(\ell+1)}(T'_i) \quad (4.41)$$

for  $0 \leq i < 2^d$ . Furthermore, we have

$$x_{\mathcal{L}-(\ell+1)}(T_i) = (x(T_i) - x(T))2^{-\mathcal{L}+(\ell+1)} \quad (4.42)$$

from which we derive

$$\begin{aligned} x_{\mathcal{L}-(\ell+1)}(T_i) &= \lambda(x(T'_i) - x(T'))2^{-\mathcal{L}+(\ell+1)} \\ &= 2^{\ell'-\ell}(x(T'_i) - x(T'))2^{-\mathcal{L}+(\ell+1)} \\ &= (x(T'_i) - x(T'))2^{-\mathcal{L}+(\ell+1)} \\ &= x_{\mathcal{L}-(\ell+1)}(T'_i), \end{aligned} \quad (4.43)$$

and analogously

$$\begin{aligned} y_{\mathcal{L}-(\ell+1)}(T_i) &= y_{\mathcal{L}-(\ell+1)}(T'_i), \\ z_{\mathcal{L}-(\ell+1)}(T_i) &= z_{\mathcal{L}-(\ell+1)}(T'_i). \end{aligned} \quad (4.44)$$

This shows that the tetrahedral Morton order of the children of  $T$  and  $T'$  are the same and  $\sigma'$  must equal  $\sigma$ .  $\square$

$I_{\text{loc}}$	Child				$I_{\text{loc}}$	Child									
	2D	$T_0$	$T_1$	$T_2$		$T_3$	$T_4$	$T_5$	$T_6$	$T_7$					
b	0	0	1	3	2	b	0	0	1	4	7	2	3	6	5
	1	0	2	3	1		1	0	1	5	7	2	3	6	4
							2	0	3	4	7	1	2	6	5
							3	0	1	6	7	2	3	4	5
							4	0	3	5	7	1	2	4	6
							5	0	3	6	7	2	1	4	5

Table 4.2.: The local index of the children of a  $d$ -simplex  $T$  according to the TM-ordering. For each type  $b$ , the  $2^d$  children  $T_0, \dots, T_{2^d-1}$  of a simplex of this type can be ordered according to their TM-indices. The position of  $T_i$  according to the TM-order is the local index  $I_{\text{loc}}(T_i) = \sigma_b(i)$ .

**Definition 4.21.** Let  $T \in \mathcal{T}_d$  such that  $T$ 's parent  $P$  has type  $b$  and  $T$  is the  $i$ th child of  $P$  according to Bey's order (4.2),  $0 \leq i < 2^d$ . We call the number  $\sigma_b(i)$  the **local index** of the  $d$ -simplex  $T$  and use the notation

$$I_{\text{loc}}(T) := \sigma_b(i) \quad (4.45)$$

to denote the child number in the TM-ordering, subsequently written TM-child. By definition, the local index of the root simplex is zero,  $I_{\text{loc}}(T_d^0) := 0$ . Table 4.2 lists the local indices for each parent type.

Thus, we know for each type  $0 \leq b < d!$  how the children of a tetrahedron of type  $b$  are traversed. This gives us an approach for describing the SFC arising from the TM-index in a recursive fashion [64]. By specifying for each possible type  $b$  the order and types of the children of a type  $b$  simplex, we can build up the SFC. In Figure 4.7 we describe the SFC for triangles in this way. In three dimensions it is not convenient to draw the six pictures for the different types, but the SFC can be derived similarly from Tables 4.1 and 4.2.

### 4.3. Low-level algorithms on simplices

In this section we present fundamental algorithms that operate on  $d$ -simplices in  $\mathcal{T}_d$ . These algorithms include computations of parent and child simplices, computation of face-neighbors and computations involved with the TM-index. To simplify the notation we carry out all algorithms for tetrahedra and then describe how to modify them for triangles. We introduce the data type **Tet** and do not distinguish between the abstract concept of a **Tet** and the geometric object (tetrahedron or triangle) that it represents. The data type **Tet**  $T$  has the following members:

- $T.\ell$  — the refinement level of  $T$ ;




Figure 4.7.: Left: Using the notation from [64] we recursively describe the SFC arising from the TM-index for triangles. The number inside each child triangle is its local index.  $R$  denotes the refinement scheme for type 0 triangles and  $F$  for type 1 triangles. This pattern can be obtained from Tables 4.1 and 4.2. Right: The SFC for an example adaptive refinement of the root triangle.

- $T.\vec{x} = (T.x, T.y, T.z)$  — the  $x$ -,  $y$ - and  $z$ -coordinates of  $T$ 's anchor node, also sometimes referred to as  $T.x_0$ ,  $T.x_1$ , and  $T.x_2$ ;
- $T.b$  — the type of  $T$ .

In 2D computations the parameter  $T.z$  is not present. To avoid confusion we use the notation  $\vec{x}$  to denote vectors in  $\mathbb{Z}^d$  and  $x$  (without arrow) for integers, thus numbers in  $\mathbb{Z}$ . From Corollary 4.9 we know that the values stored in a **Tet** suffice to uniquely identify a  $d$ -simplex  $T \in \mathcal{T}$ .

**Remark 4.22** (Storage requirement). The algorithms that we present in this section only need this data as input for a simplex resulting in a fixed storage size per **Tet**. If, for example, the maximum level  $\mathcal{L}$  is 32 or less, then the coordinates can be stored in one 4-byte integer per dimension, while the level and type occupy one byte each, leading to a total storage of

$$\begin{aligned} 2 \times 4 + 1 + 1 &= 10 \text{ bytes per Tet in 2D,} \\ 3 \times 4 + 1 + 1 &= 14 \text{ bytes per Tet in 3D.} \end{aligned} \tag{4.46}$$

**Remark 4.23** (Runtime). Most of these algorithms run in constant time independent of the maximum level  $\mathcal{L}$ . The only operations using a loop over the level  $\mathcal{L}$  or  $T.l$ , thus having  $\mathcal{O}(\mathcal{L})$  runtime, are computing the consecutive index from a **Tet** and initializing a **Tet** according to a given consecutive index. Hence, we show how to replace repetitive calls of these relatively involved algorithms by more efficient constant-time ones.

### 4.3.1. The coordinates of a $d$ -simplex

The coordinates of the  $d + 1$  nodes of a  $d$ -simplex  $T$  can be obtained easily from its Tet-id, the relation (4.3), and simple arithmetic: If  $T$  is a  $d$ -simplex of level  $\ell$ , type

$b$  and anchor node  $\vec{x}_0 \in \mathbb{Z}^d$ , then

$$T = 2^{\mathcal{L}-\ell} S_b + \vec{x}_0. \quad (4.47)$$

Hence, in order to compute the coordinates of  $T$  we can take the coordinates of  $S_b$ , as given in (4.3), and then use relation (4.47). A closer look at (4.3) reveals that it is not necessary to examine all coordinates of  $S_b$  in order to compute the  $x_i$ , but that they can also be computed arithmetically. This computation is carried out in Algorithm 4.3.1.

---

**Algorithm 4.3.1:** Coordinates(Tet  $T$ )

---

**Result:** Array of coordinates of all of  $T$ 's vertices.

```

1  $X \leftarrow (T.\vec{x}, 0, 0, 0)$ 
2  $h \leftarrow 2^{\mathcal{L}-\ell}$ 
3  $i \leftarrow \lfloor \frac{T.b}{2} \rfloor$  /* Replace with  $i \leftarrow T.b$  for 2D */
4 if  $T.b \% 2 = 0$  then
5   |  $j \leftarrow (i + 2) \% 3$ 
6 else
7   |  $j \leftarrow (i + 1) \% 3$ 
8  $X[1] \leftarrow X[0] + he_i$ 
9  $X[2] \leftarrow X[1] + he_j$  /* Replace with  $X[2] \leftarrow X[0] + (h, h)^T$  for 2D */
10  $X[3] \leftarrow X[0] + (h, h, h)^T$  /* Remove this line for 2D */
11 return  $X$ 
```

---

### 4.3.2. Parent and child

In this section we describe how to compute the Tet-ids of the parent  $P(T)$  and of the  $2^d$  children  $T_i$ ,  $0 \leq i < 2^d$ , of a given  $d$ -simplex  $T \in \mathcal{T}_d$ . Computing the anchor node coordinates of the parent is easy, since their first  $T.\ell - 1$  bits correspond to the coordinates of  $T$ 's anchor node and the rest of their bits is zero. For computing the type of  $P(T)$ , we need the function

$$\begin{aligned} \text{Pt: } \{0, \dots, 2^d - 1\} \times \{0, \dots, d! - 1\} &\longrightarrow \{0, \dots, d! - 1\}, \\ (\text{cube-id}(T), T.b) &\longmapsto P.b, \end{aligned} \quad (4.48)$$

giving the type of  $T$ 's parent in dependence of its cube-id and type. In Figure 4.8 we list all values of this function for  $d \in \{2, 3\}$ .

The algorithm **Parent** to compute the parent of  $T$  now puts these two ideas together, computing the coordinates and type of  $P(T)$ . Algorithm 4.3.3 shows an implementation. It uses Algorithm 4.3.2 to compute the cube-id of a  $d$ -simplex.

For the computation of one child  $T_i$  of  $T$  for a given  $i \in \{0, \dots, 2^d - 1\}$  we look at Bey's definition of the subsimplices in (4.2) and see that in order to compute the anchor node of the child we need to know some of the node coordinates  $\vec{x}_0, \dots, \vec{x}_d$  of the parent simplex  $T$ . These can be obtained via Algorithm 4.3.1. However, it is more efficient to compute only those coordinates of  $T$  that are actually necessary.

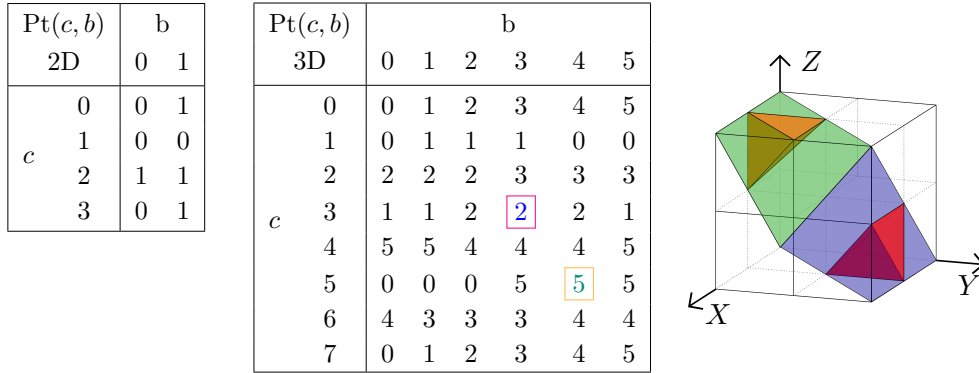


Figure 4.8.: The type of the parent of a  $d$ -simplex  $T$  can be determined from  $T$ 's cube-id  $c$  and type  $b$ . Left: The values of  $Pt$  from (4.48) in the 2D case. Middle: The values of the function  $Pt$  in the 3D case. Right: Two examples showing the computation in 3D. (1) The small tetrahedron in the top left corner (orange) has cube-id 5 and type 4, and its parent (green) can be seen to have type 5. (2) The small tetrahedron at the bottom right (red) has cube-id 3 and type 3, and its parent (blue) has type 2.

---

**Algorithm 4.3.2:**  $c\text{-id}(\text{Tet } T, \text{int } \ell)$

---

**Result:** The cube-id of  $T$ .

```

1  $i \leftarrow 0, h \leftarrow 2^{\mathcal{L}-\ell}$ 
2  $i \mid= (T.x \ \& \ h) ? 1 : 0$ 
3  $i \mid= (T.y \ \& \ h) ? 2 : 0$ 
4  $i \mid= (T.z \ \& \ h) ? 4 : 0$ 
5 return  $i$ 

```

/\* Remove this line for 2D \*/

---



---

**Algorithm 4.3.3:**  $\text{Parent}(\text{Tet } T)$

---

**Result:** The parent  $P$  of  $T$ .

```

1  $h \leftarrow 2^{\mathcal{L}-T.\ell}$ 
2  $P.\ell \leftarrow T.\ell - 1$ 
3  $P.x \leftarrow T.x \ \& \ \neg h$ 
4  $P.y \leftarrow T.y \ \& \ \neg h$ 
5  $P.z \leftarrow T.z \ \& \ \neg h$ 
6  $P.b \leftarrow Pt(c\text{-id}(T, T.\ell), T.b)$ 
7 return  $P$ 

```

/\* Remove this line for 2D \*/  
/\* See (4.48) and Figure 4.8 for Pt \*/

---

To compute the Tet-id of  $T_i$  we also need to know its type. The type of  $T_i$  depends only on the type of  $T$ , and in the algorithm we use the function **Ct** (children type) to compute this type. **Ct** is effectively an evaluation of Table 4.1. Algorithm 4.3.4 shows now how to compute the coordinates of the  $i$ th child of  $T$  in Bey’s order.

When we would like to compute the  $i$ th child of a  $d$ -simplex  $T$  of type  $b$  with respect to the tetrahedral Morton order (thus the child  $T_k$  of  $T$  with  $I_{\text{loc}}(T_k) = i$ ) we just call Algorithm 4.3.4 with  $\sigma_b^{-1}(i)$  as input. The permutations  $\sigma_b^{-1}$  are available from Table 4.2; see (4.45) and Algorithm 4.3.5.

---

**Algorithm 4.3.4:** Child(Tet  $T$ ,int  $i$ )

---

**Result:** The  $i$ -th child (in Bey’s order)  $T_i$  of  $T$ .

```

1  $X \leftarrow \text{Coordinates}(T)$ 
2 if  $i = 0$  then  $j \leftarrow 0$ 
3 else if  $i \in \{1, 4, 5\}$  then  $j \leftarrow 1$ 
4 else if  $i \in \{2, 6, 7\}$  then  $j \leftarrow 2$ 
5 else if  $i = 3$  then  $j \leftarrow 3$                                /* If  $i = 3$  then  $j \leftarrow 1$  for 2D */
6  $T_i.\vec{x} \leftarrow \frac{1}{2}(X[0] + X[j])$ 
7  $T_i.b \leftarrow \text{Ct}(T.b, i)$                                    /* See Table 4.1 */
```

---



---

**Algorithm 4.3.5:** TM-Child(Tet  $T$ ,int  $i$ )

---

**Result:** The  $i$ -th child of  $T$  in TM-order.

```

1 return Child( $T, \sigma_{T.b}^{-1}(i)$ )                               /* See Table 4.2 */
```

---

### 4.3.3. Neighbor simplices

Many applications—e.g., finite element methods—need to gather information about the face-neighboring simplices of a given simplex in a refinement. In this section we describe a level-independent constant-runtime algorithm to compute the Tet-id of the same level neighbor along a specific face  $f$  of a given  $d$ -simplex  $T$ . This algorithm is very lightweight since it only requires a few arithmetic computations involving the Tet-id of  $T$  and the number  $f$ . In comparison to other approaches to finding neighbors in constant time [21, 89], our algorithm does not involve the computation of any of  $T$ ’s ancestors.

The  $d + 1$  faces of a  $d$ -simplex  $T = [\vec{x}_0, \dots, \vec{x}_d]$  are numbered  $f_0, \dots, f_d$  in such a way that face  $f_i$  is the face not containing the node  $\vec{x}_i$ . To examine the situation where two  $d$ -simplices of the same level share a common face, let  $\mathcal{T}_d^\ell$  denote a uniform refinement of  $T_d^0$  of a given level  $0 \leq \ell \leq \mathcal{L}$ ,

$$\mathcal{T}_d^\ell := \{ T \mid T \text{ is a descendant of } T_d^0 \text{ of level } \ell \} \subset \mathcal{T}_d. \quad (4.49)$$

$\mathcal{T}_d$  can be seen as the disjoint union of all the  $\mathcal{T}_d^\ell$ :

$$\mathcal{T}_d = \bigcup_{\ell=0}^{\mathcal{L}} \mathcal{T}_d^\ell. \quad (4.50)$$




Figure 4.9.: A tetrahedron  $T$  of type 5 (in the middle, red) and its four face-neighbors (blue) of types 1, 0, 4, and 3, drawn with their associated cubes (exploded view). We see that the type of  $T$ 's neighbors depends only on its type, while their node coordinates relative to  $T$ 's depend additionally on  $T$ 's level.

Given a  $d$ -simplex  $T \in \mathcal{T}_d^\ell$  and a face number  $i \in \{0, \dots, d\}$ , denote  $T$ 's neighbor in  $\mathcal{T}_d^\ell$  across face  $f = f_i$  by  $\mathcal{N}_f(T)$ , and denote the face number of the neighbor simplex  $\mathcal{N}_f(T)$  across which  $T$  is its neighbor by  $\tilde{f}(T)$ . Hence, the relation

$$\mathcal{N}_{\tilde{f}(T)}(\mathcal{N}_f(T)) = T \quad (4.51)$$

holds for each face  $f$  of  $T$ .

Our aim is to compute the Tet-id of  $\mathcal{N}_f(T)$  and  $\tilde{f}(T)$  from the Tet-id of  $T$ . Using the underlying cube structure this problem can be solved for each occurring type of  $d$ -simplex separately, and the solution scheme is independent of the coordinates of  $T$  and of  $\ell$ . In Figure 4.9 the situation for a tetrahedron of type 5 is illustrated, and in Tables 4.3 and 4.4 we present the general solution for each type.

Using these tables, a constant-time computation of the Tet-id of  $\mathcal{N}_f(T)$  and of  $\tilde{f}(T)$  from the Tet-id of  $T$  is possible, and the 3D case is carried out in Algorithm 4.3.6. Note that this algorithm uses arithmetic expressions in  $T.b$  to avoid the sixfold distinction of cases.

**Remark 4.24.** To find existing neighbors in a nonuniform refinement we use Algorithm 4.3.6 in combination with `Parent` and `TM-Child` and comparison functions.

Of course it is possible that  $\mathcal{N}_f(T)$  does not belong to  $\mathcal{T}_d^\ell$  any longer. If this is the case, then  $f$  was part of the boundary of the root simplex  $T_d^0$ . We describe in the next section how we can decide in constant time whether  $\mathcal{N}_f(T)$  is in  $T_d^\ell$  or not.

For completeness, we summarize the geometry and numbers of  $d$ -simplices touching each other via a corner ( $d = 2$  or  $d = 3$ ) or edge (only  $d = 3$ ). In this chapter we do not list these neighboring tetrahedra in detail.

For  $d = 3$  each corner in the mesh  $\mathcal{T}_3^\ell$  has 24 adjacent tetrahedra; thus each tetrahedron has at each corner 23 other tetrahedra that share this particular corner. For  $d = 2$  the situation is similar, with six triangles meeting at each corner. To examine the number of adjacent tetrahedra to an edge we distinguish three types of edges in  $\mathcal{T}_d^\ell$ :



2D	$f$	0	1	2
$T.b = 0$	$\mathcal{N}_f.b$	1	1	1
	$\mathcal{N}_f.x$	$T.x + h$	$T.x$	$T.x$
	$\mathcal{N}_f.y$	$T.y$	$T.y$	$T.y - h$
	$\tilde{f}$	2	1	0
$T.b = 1$	$\mathcal{N}_f.b$	0	0	0
	$\mathcal{N}_f.x$	$T.x$	$T.x$	$T.x - h$
	$\mathcal{N}_f.y$	$T.y + h$	$T.y$	$T.y$
	$\tilde{f}$	2	1	0

Table 4.3.: Face neighbors in 2D. For each possible type  $b \in \{0, 1\}$  of a triangle  $T$  and each of its faces  $f = f_i$ ,  $i \in \{0, 1, 2\}$ , the type, anchor node coordinates, and corresponding face number  $\tilde{f}$  of  $T$ 's neighbor across  $f$  are shown. In the 2D case we can directly compute  $\mathcal{N}.b = 1 - T.b$  and  $\tilde{f} = 2 - f$ . Here,  $h = 2^{\mathcal{L}-\ell}$  refers to the length of one square of level  $\ell$ .

3D	$f$	0	1	2	3	3D	$f$	0	1	2	3
$T.b = 0$	$\mathcal{N}_f.b$	4	5	1	2	$T.b = 3$	$\mathcal{N}_f.b$	5	4	2	1
	$\mathcal{N}_f.x$	$T.x + h$	$T.x$	$T.x$	$T.x$		$\mathcal{N}_f.x$	$T.x$	$T.x$	$T.x$	$T.x - h$
	$\mathcal{N}_f.y$	$T.y$	$T.y$	$T.y$	$T.y - h$		$\mathcal{N}_f.y$	$T.y + h$	$T.y$	$T.y$	$T.y$
	$\mathcal{N}_f.z$	$T.z$	$T.z$	$T.z$	$T.z$		$\mathcal{N}_f.z$	$T.z$	$T.z$	$T.z$	$T.z$
	$\tilde{f}$	3	1	2	0		$\tilde{f}$	3	1	2	0
$T.b = 1$	$\mathcal{N}_f.b$	3	2	0	5	$T.b = 4$	$\mathcal{N}_f.b$	2	3	5	0
	$\mathcal{N}_f.x$	$T.x + h$	$T.x$	$T.x$	$T.x$		$\mathcal{N}_f.x$	$T.x$	$T.x$	$T.x$	$T.x - h$
	$\mathcal{N}_f.y$	$T.y$	$T.y$	$T.y$	$T.y$		$\mathcal{N}_f.y$	$T.y$	$T.y$	$T.y$	$T.y$
	$\mathcal{N}_f.z$	$T.z$	$T.z$	$T.z$	$T.z - h$		$\mathcal{N}_f.z$	$T.z + h$	$T.z$	$T.z$	$T.z$
	$\tilde{f}$	3	1	2	0		$\tilde{f}$	3	1	2	0
$T.b = 2$	$\mathcal{N}_f.b$	0	1	3	4	$T.b = 5$	$\mathcal{N}_f.b$	1	0	4	3
	$\mathcal{N}_f.x$	$T.x$	$T.x$	$T.x$	$T.x$		$\mathcal{N}_f.x$	$T.x$	$T.x$	$T.x$	$T.x$
	$\mathcal{N}_f.y$	$T.y + h$	$T.y$	$T.y$	$T.y$		$\mathcal{N}_f.y$	$T.y$	$T.y$	$T.y$	$T.y - h$
	$\mathcal{N}_f.z$	$T.z$	$T.z$	$T.z$	$T.z - h$		$\mathcal{N}_f.z$	$T.z + h$	$T.z$	$T.z$	$T.z$
	$\tilde{f}$	3	1	2	0		$\tilde{f}$	3	1	2	0

Table 4.4.: Face neighbors in 3D. For each possible type  $b \in \{0, 1, 2, 3, 4, 5\}$  of a tetrahedron  $T$  and each of its faces  $f = f_i$ ,  $i \in \{0, 1, 2, 3\}$  the type  $\mathcal{N}_f(T).b$  of  $T$ 's neighbor across  $f$ , its coordinates of the anchor node  $\mathcal{N}_f(T).x$ ,  $\mathcal{N}_f(T).y$ ,  $\mathcal{N}_f(T).z$  and the corresponding face number  $\tilde{f}(T)$ , across which  $T$  is  $\mathcal{N}_f(T)$ 's neighbor, are shown.

---

**Algorithm 4.3.6: Face-neighbor(Tet  $T$ , int  $f$ )**

---

```
Result: The face-neighbor  $\mathcal{N}_f(T)$  of  $T$  across  $f$  and the corresponding face  $\tilde{f}$ .
1  $b \leftarrow T.b, x_0 \leftarrow T.x_0, x_1 \leftarrow T.x_1, x_2 \leftarrow T.x_2$ 
2 if  $f = 1$  or  $f = 2$  then
3    $\tilde{f} \leftarrow f$  if  $(b \% 2 = 0$  and  $f = 2)$  or  $(b \% 2 \neq 0$  and  $f = 1)$  then
4      $b \leftarrow b + 1$ 
5   else
6      $b \leftarrow b - 1$ 
7 else
8    $\tilde{f} \leftarrow 3 - f$ 
9    $h \leftarrow 2^{\mathcal{L}-T.\ell}$ 
10  if  $f = 0$  then /*  $f = 0$  */
11     $i \leftarrow b \text{ div } 2$ 
12     $x_i \leftarrow T.x_i + h$ 
13     $b \leftarrow b + (b \% 2 = 1 ? 2 : 4)$ 
14  else /*  $f = 3$  */
15     $i \leftarrow (b + 3) \% 6 \text{ div } 2$ 
16     $x_i \leftarrow T.x_i - h$ 
17     $b \leftarrow b + (b \% 2 = 0 ? 2 : 4)$ 
18  $N.\vec{x} \leftarrow (x_0, x_1, x_2)$ 
19  $N.\ell \leftarrow T.\ell$ 
20  $N.b \leftarrow b \% 6$ 
21 return  $(N, \tilde{f})$ 
```

---

1. edges that are also edges in the underlying hexahedral mesh;
2. edges that are the diagonal of a side of a cube in the hexahedral mesh;
3. edges that correspond to the inner diagonal of a cube in the hexahedral mesh.

Edges of the first and third kind have six adjacent tetrahedra each, and edges of the second kind do have four adjacent tetrahedra each.

#### 4.3.4. The exterior of the root simplex

When computing neighboring  $d$ -simplices it is possible that the neighbor in question does not belong to the root simplex  $T_d^0$  but lies outside of it. If we look at face-neighbors of a  $d$ -simplex  $T$ , the fact that the considered neighbor lies outside means that the respective face was on the boundary of  $T_d^0$ . In order to check whether a computed  $d$ -simplex is outside the base simplex, we investigate a more general problem: Given anchor node coordinates  $(x_0, y_0)^T \in \mathbb{Z}^2$ , respectively  $(x_0, y_0, z_0)^T \in \mathbb{Z}^3$ , of type  $b$  a level  $\ell$ , decide whether the corresponding  $d$ -simplex  $N$  lies inside or outside of the root tetrahedron  $T_d^0$ :  $N \in \mathcal{T}_d^\ell$  or  $N \notin \mathcal{T}_d^\ell$ . At the end of this section we furthermore generalize this to the problem of deciding for any two  $d$ -simplices  $N$  and  $T$  whether or not  $N$  lies outside of  $T$ . We solve this problem in constant time and independent of the levels of  $N$  and  $T$ .

We examine the 3D case. Looking at  $T_3^0$  we observe that two of its boundary faces correspond to faces of the root cube, namely, the intersections of  $T_3^0$  with the  $y = 0$  and the  $x = 2^{\mathcal{L}}$  planes. The other two boundary faces of  $T_3^0$  are the intersections with the  $x = z$  and the  $y = z$  planes. Thus, the boundary of  $T_3^0$  can be described as the intersection of  $T_3^0$  with those planes. We refer to the latter two planes as  $E_1$  and  $E_2$ .

Let  $N$  be a tetrahedron with anchor node  $(x_0, y_0, z_0)^T \in \mathbb{Z}^3$  of type  $b$  and level  $\ell$  and denote with  $(x_i, y_i, z_i)^T$  the coordinates of node  $i$  of  $N$ . Since  $(x_i, y_i, z_i)^T \geq (x_0, y_0, z_0)^T$  (componentwise), we directly conclude that if  $x_0 \geq 2^{\mathcal{L}}$  or  $y_0 < 0$  then  $N \notin \mathcal{T}_3$ . Because the outer normal vectors of  $T_3^0$  on the two faces intersecting with  $E_1$  and  $E_2$  are

$$\vec{n}_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \quad \text{and} \quad \vec{n}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}, \quad (4.52)$$

we also conclude that  $N \notin \mathcal{T}_3$  if  $z_0 - x_0 > 0$  or  $y_0 - z_0 > 0$ . Now we have already covered all the cases except those where the anchor node of  $N$  lies directly in  $E_1$  or  $E_2$ . In these cases we cannot solve the problem by looking at the coordinates of the anchor node alone, since there exist tetrahedra  $T' \in \mathcal{T}_3$  with anchor nodes lying in one of these planes (see Figure 4.10 for an illustration of the analogous case in 2D). This depends on the type of the tetrahedron in question. We observe that a tetrahedron  $T' \in \mathcal{T}_3$  with anchor node lying in  $E_1$  can have the types 0, 1, or 2, and a tetrahedron with anchor node lying in  $E_2$  can have the types 0, 4, or 5. We conclude that to check whether  $N$  is outside of the root tetrahedron we have to check if any one of six conditions is fulfilled. In fact these conditions fit into the general form below with  $x_i = x$ ,  $x_j = y$ ,  $x_k = z$ , and  $T$  as the root tetrahedron; thus  $T.x = T.y = T.z = 0$  and  $T.b = 0$ .

These generalized conditions solve the problem to check for any two given tetrahedra  $N$  and  $T$ , whether  $N$  lies outside of  $T$  or not.

**Proposition 4.25.** *Given two  $d$ -simplices  $N, T$  with  $N.\ell \geq T.\ell$ , then  $N$  is outside of the simplex  $T$ —which is equivalent to saying that  $N$  is no descendant of  $T$ —if and only if at least one of the following conditions is fulfilled.*

For 2D,

$$N.x_i - T.x_i \geq 2^{T.\ell}, \quad (4.53a)$$

$$N.x_j - T.x_j < 0, \quad (4.53b)$$

$$(N.x_j - T.x_j) - (N.x_i - T.x_i) > 0, \quad (4.53c)$$

$$N.x_i - T.x_i = N.x_j - T.x_j \quad \text{and} \quad N.b = \begin{cases} 1 & \text{if } T.b = 0, \\ 0 & \text{if } T.b = 1. \end{cases} \quad (4.53d)$$

2D	$T.b$		3D	$T.b$					
	0	1		0	1	2	3	4	5
$x_i$	$x$	$y$	$x_i$	$x$	$x$	$y$	$y$	$z$	$z$
$x_j$	$y$	$x$	$x_j$	$y$	$z$	$z$	$x$	$x$	$y$
			$x_k$	$z$	$y$	$x$	$z$	$y$	$x$

Table 4.5.: Following the general scheme described in this section to compute whether a given  $d$ -simplex  $N$  lies outside of another given  $d$ -simplex  $T$ , we give the coordinates  $x_i$ ,  $x_j$ , and  $x_k$  in dependence of the type of  $T$ .

For 3D,

$$N.x_i - T.x_i \geq 2^{\mathcal{L}-T.\ell}, \quad (4.54a)$$

$$N.x_j - T.x_j < 0, \quad (4.54b)$$

$$(N.x_k - T.x_k) - (N.x_i - T.x_i) > 0, \quad (4.54c)$$

$$(N.x_j - T.x_j) - (N.x_k - T.x_k) > 0, \quad (4.54d)$$

$$N.x_j - T.x_j = N.x_k - T.x_k$$

and  $N.b \in \begin{cases} \{T.b + 1, T.b + 2, T.b + 3\}, & \text{if } T.b \text{ is even,} \\ \{T.b - 1, T.b - 2, T.b - 3\}, & \text{if } T.b \text{ is odd,} \end{cases} \quad (4.54e)$

$$N.x_k - T.x_k = N.x_i - T.x_i$$

and  $N.b \in \begin{cases} \{T.b - 1, T.b - 2, T.b - 3\}, & \text{if } T.b \text{ is even,} \\ \{T.b + 1, T.b + 2, T.b + 3\}, & \text{if } T.b \text{ is odd.} \end{cases} \quad (4.54f)$

$$N.x_j - T.x_j = N.x_k - T.x_k \quad \text{and} \quad N.x_i - T.x_i = N.x_k - T.x_k \quad \text{and} \quad N.b \neq T.b \quad (4.54g)$$

The coordinates  $x_i$ ,  $x_j$ , and  $x_k$  are chosen in dependence of the type of  $T$  according to Table 4.5.

*Proof.* By shifting  $N$  by  $(-T.\vec{x})$  we reduce the problem to checking whether the shifted  $d$ -simplex lies outside of a  $d$ -simplex with anchor node  $\vec{0}$ , level  $T.\ell$  and type  $T.b$ . For  $d = 3$  the proof is analogous to the above argument, where we considered the case  $b = 0$  and  $\ell = 0$ . In two dimensions the situation is even simpler, since there exists only one face of the root triangle that is not a coordinate axis (see Figure 4.10).  $\square$

### 4.3.5. A consecutive index for uniform refinements

In contrast to the Morton index for cubes, the TM-index for  $d$ -simplices does not produce a consecutive range of numbers. Therefore, two simplices  $T$  and  $T'$  of level  $\ell$




Figure 4.10.: A uniform level 2 refinement of the triangle  $T^0$  in 2D and triangles lying outside of it with their anchor nodes marked. When deciding whether a triangle with given anchor node coordinates lies outside of  $T^0$  there are four cases to consider, one for each face of  $T^0$ . For the two faces lying parallel to the  $X$ -axis, respectively,  $Y$ -axis, it suffices to check whether the  $x$ -coordinate is greater than or equal to  $2^{\mathcal{L}}$ , or whether the  $y$ -coordinate is smaller than 0. Similarly one can conclude that the triangle lies outside of  $T^0$  if its  $x$ -coordinate is smaller than its  $y$ -coordinate. If both coordinates agree (and none of the previous cases applies) then the given triangle is outside  $T^0$  if and only if its type is 1.

that are direct successors/predecessors with respect to the tetrahedral Morton order do not necessarily fulfill  $m(T) = m(T') \pm 2^{d(\mathcal{L}-\ell)}$  or  $m(T) = m(T') \pm 1$ . For  $d = 3$  this follows directly from the fact that each  $b^j$  occupies three bits, but there are only six values that each  $b^j$  can assume, since there are only six different types. In 2D this follows from the fact that not every combination of anchor node coordinates and type can occur for triangles in  $\mathcal{T}_2$ , the triangle with anchor node  $(0, 0)$  and type 1 being one example. This also means that the largest occurring TM-index is bigger than  $2^{d\mathcal{L}} - 1$ .

Constructing a consecutive index as in Definition 3.16 that respects the order given by the TM-index is possible, as we show in this section. Since in practice it is more convenient to work with this consecutive index instead of the TM-index, our aim is to construct for each uniform refinement level  $\ell$  a consecutive index  $\mathcal{I}_\ell(T) \in \{0, \dots, 2^{d\ell} - 1\}$ , such that

$$\forall T, S \in \mathcal{T}_d^\ell: \quad \mathcal{I}_\ell(T) < \mathcal{I}_\ell(S) \Leftrightarrow m(T) < m(S). \quad (4.55)$$

This index can also be understood as a bijection

$$\mathcal{I}_\ell: \left\{ m \in \mathbb{N}_0 \mid m = m(T) \text{ for a } T \in \mathcal{T}_d^\ell \right\} \xrightarrow{\cong} \left\{ 0, \dots, 2^{d\ell} - 1 \right\}, \quad (4.56)$$

mapping the TM-indices of level  $\ell$   $d$ -simplices to a consecutive range of numbers. See also Definition 3.16. It is obvious that  $\mathcal{I}_\ell(T_d^0) = 0$ . The index  $\mathcal{I}_\ell(T)$  can be easily computed as the  $\ell$ -digit  $2^d$ -ary number consisting of the local indices as digits, thus

$$\mathcal{I}_\ell(T) = (I_{\text{loc}}(T^1), \dots, I_{\text{loc}}(T^\ell))_{2^d}. \quad (4.57)$$

Algorithm 4.3.7 shows an implementation of this computation. It can be done directly from the Tet-id of  $T$ , and thus it is not necessary to compute the TM-index of  $T$  first.

---

**Algorithm 4.3.7:**  $\mathcal{I}(\text{Tet}T)$

---

**Result:** The consecutive index  $\mathcal{I}_{T,\ell}(T)$ .

```

1  $I \leftarrow 0, b \leftarrow T.b$ 
2 for  $T.\ell \geq i \geq 1$  do
3    $c \leftarrow \text{c-id}(T, i)$ 
4    $I \leftarrow I + 8^i I_{\text{loc}}^b(c)$            /* See Table 4.6; multiply with  $4^i$  for 2D */
5    $b \leftarrow \text{Pt}(c, b)$ 
6 return  $I$ 
```

---

The inverse operation of computing  $T$  from  $\mathcal{I}_\ell(T)$  and a given level  $\ell$  can be carried out in a similar fashion; see Algorithm 4.3.8. For each  $0 \leq i \leq \ell$  we look up the type  $b$  and the cube-id of  $T^i$  from  $I_{\text{loc}}(T^i)$  and the type of  $\text{Parent}(T^i) = T^{i-1}$  (starting with  $\text{type}(T^0) = 0$ ) via Tables 4.7 and 4.8. From the cube-ids we can build up the anchor node coordinates of  $T$ . The last computed type is the type of  $T$ . The runtime of this algorithm is  $\mathcal{O}(T.\ell)$ .

---

**Algorithm 4.3.8:** T(consecutive index  $\mathcal{I}$ ,int  $\ell$ )

---

**Result:** The simplex  $T$  with  $\mathcal{I}_\ell(T) = \mathcal{I}$ .

```
1  $T.x, T.y, T.z \leftarrow 0, b \leftarrow 0$ 
2 for  $1 \leq i \leq \ell$  do
3   Get  $I_{\text{loc}}(T^i)$  from  $\mathcal{I}$  /* See (4.57) */
4    $c \leftarrow c\text{-id}(T^i), b \leftarrow T^i.b$  /* See Tables 4.7 and 4.8 */
5   if  $c \& 1$  then  $T.x \leftarrow T.x + 2^{\mathcal{L}-i}$ 
6   if  $c \& 2$  then  $T.y \leftarrow T.y + 2^{\mathcal{L}-i}$ 
7   if  $c \& 4$  then  $T.z \leftarrow T.z + 2^{\mathcal{L}-i}$  /* Remove this line for 2D */
8  $T.b \leftarrow b$ 
9 return  $T$ 
```

---

Similar to Algorithm 4.3.7 is Algorithm `Is_valid`, which decides whether a given index  $m \in [0, 2^{6\mathcal{L}}) \cap \mathbb{Z}$  is in fact a TM-index for a tetrahedron. Thus, in the spirit of Section 4.2.5 we can decide whether a given 6D cube is in the image of map (4.27) that embeds  $\mathcal{T}_3$  into the set of 6D subcubes of  $[0, 2^\mathcal{L}]^6$ . The runtime of `Is_valid` is  $\mathcal{O}(\mathcal{L})$ .

---

**Algorithm 4.3.9:** `Is_valid`( $m \in [0, 2^{6\mathcal{L}}) \cap \mathbb{Z}, \ell$ )

---

**Result:** True if there exists a simplex  $T$  with  $m(T) = m$ , false otherwise.

```
1  $k \leftarrow 6(\mathcal{L} - i)$ 
2 for  $\ell \geq i \geq 1$  do
3    $b \leftarrow (m_k, m_{k+1}, m_{k+2})_8$ 
4    $c \leftarrow (m_{k+3}, m_{k+4}, m_{k+5})_8$ 
5    $k \leftarrow 6(\mathcal{L} - i + 1)$ 
6   if  $(m_k, m_{k+1}, m_{k+2})_8 \neq Pt(c, b)$  then /* Take  $(0, 0, 0)_8$  if  $i = 1$  */
7     return False
8 return True
```

---

The consecutive index simplifies the relation between the TM-index of a simplex and its position in the SFC. In the special case of a uniform mesh, the consecutive index and the position are identical.

### 4.3.6. Successor and predecessor

Calculating the TM-index corresponding to a particular consecutive index is occasionally needed in high-level algorithms. This is relatively expensive, since it involves a loop over all refinement levels, thus some 10 to 30 in extreme cases. However often the task is to compute a whole range of  $d$ -simplices. This occurs, for example, when creating an initial uniform refinement of a given mesh (see Algorithm `New` in Section 4.4.1). That is, for a given consecutive index  $\mathcal{I}$ , a level  $\ell$ , and a count  $n$ , find the  $n$  level- $\ell$  simplices following the  $d$ -simplex corresponding to the consecutive index  $\mathcal{I}$ , that is, the  $d$ -simplices corresponding to the  $n$  consecutive indices  $\mathcal{I}, \mathcal{I} + 1, \dots, \mathcal{I} + n - 1$ . Ideally, this operation should run linearly in  $n$ , in-

$I_{\text{loc}}^b(c)$	cube-id c				$I_{\text{loc}}^b(c)$	cube-id c									
2D	0	1	2	3	3D	0	1	2	3	4	5	6	7		
b	0	0	1	1	3	b	0	0	1	1	4	1	4	4	7
	1	0	2	2	3		1	0	1	2	5	2	5	4	7
							2	0	2	3	4	1	6	5	7
							3	0	3	1	5	2	4	6	7
							4	0	2	2	6	3	5	5	7
							5	0	3	3	6	3	6	6	7

Table 4.6.: The local index of a tetrahedron  $T \in \mathcal{T}$  in dependence of its cube-id  $c$  and type  $b$ .

cube-id( $T$ )	$I_{\text{loc}}(T)$				cube-id( $T$ )	$I_{\text{loc}}(T)$									
2D	0	1	2	3	3D	0	1	2	3	4	5	6	7		
$P.b$	0	0	1	1	3	$P.b$	0	0	1	1	5	5	5	7	
	1	0	2	2	3		1	0	1	1	1	3	3	3	7
							2	0	2	2	2	3	3	3	7
							3	0	2	2	2	6	6	6	7
							4	0	4	4	4	6	6	6	7
							5	0	4	4	4	5	5	5	7

Table 4.7.: For a tetrahedron  $T \in \mathcal{T}$  of local index  $I_{\text{loc}}$  whose parent  $P$  has type  $P.b$  we give the cube-id of  $T$ .

$T.b$	$I_{\text{loc}}(T)$				$T.b$	$I_{\text{loc}}(T)$									
2D	0	1	2	3	3D	0	1	2	3	4	5	6	7		
$P.b$	0	0	0	1	0	$P.b$	0	0	0	4	5	0	1	2	0
	1	1	0	1	1		1	1	1	2	3	0	1	5	1
							2	2	0	1	2	2	3	4	2
							3	3	3	4	5	1	2	3	3
							4	4	2	3	4	0	4	5	4
							5	5	0	1	5	3	4	5	5

Table 4.8.: For a tetrahedron  $T \in \mathcal{T}$  of local index  $I_{\text{loc}}$  whose parent  $P$  has type  $P.b$  we give the type of  $T$ .



dependent of  $\ell$ , but if we used Algorithm 4.3.8 to create each of the  $n + 1$  simplices we would have a runtime of  $\mathcal{O}(n\mathcal{L})$ . In order to achieve the desired linear runtime we introduce the operations **Successor** and **Predecessor** that run in average  $\mathcal{O}(1)$  time. These operations compute from a given  $d$ -simplex  $T$  of level  $\ell$  with consecutive index  $\mathcal{I}_\ell(T)$  the  $d$ -simplex  $T'$  whose consecutive index is  $\mathcal{I}_\ell(T) + 1$ , respectively,  $\mathcal{I}_\ell(T) - 1$ . Thus,  $T'$  is the next level  $\ell$  simplex in the SFC after  $T$  (resp. the previous one). Algorithm 4.3.10, which we introduce to solve this problem does not require knowledge about the consecutive indices  $\mathcal{I}_\ell(T)$  and  $\mathcal{I}_\ell(T) \pm 1$  and can be computed significantly faster than Algorithm 4.3.8; see Lemma 4.26.

---

**Algorithm 4.3.10: Successor(Tet  $T$ )**

---

**Result:** The successor  $T'$  of  $T$ .

```

1 return Successor_recursion( $T, T, T.\ell$ )

Function Successor_recursion(Tet  $T, \text{Tet } T', \text{int } \ell$ )
1  $c \leftarrow c\text{-id}(T, \ell)$ 
2 From  $c$  and  $b$  look up  $i := I_{\text{loc}}(T^\ell)$  /* See Table 4.6 */
3  $i \leftarrow (i + 1) \% 8$ 
4 if  $i = 0$  then /* Enter recursion (in rare cases) */
5 |    $T' \leftarrow \text{Successor\_recursion}(T, T', \ell - 1)$ 
6 |    $\hat{b} \leftarrow T'.b$  /*  $\hat{b}$  stores the type of  $T'^{\ell-1}$  */
7 else
8 |    $\hat{b} \leftarrow \text{Pt}(c, b)$ 
9 From  $\hat{b}$  and  $I_{\text{loc}} = i$  look up  $(c', b')$  /* See Tables 4.7 and 4.8 */
10 Set the level  $\ell$  entries of  $T'.x, T'.y$  and  $T'.z$  to  $c'$ 
11  $T'.b \leftarrow b'$ 
12 return  $T'$ 
```

---

To compute the predecessor of  $T$  we only need to reverse the sign in Line 3 in the **Successor\_recursion** subroutine of Algorithm 4.3.10.

**Lemma 4.26.** *Algorithm 4.3.10 has constant average runtime (independent of  $\mathcal{L}$ ).*

*Proof.* Because each operation in the algorithm can be executed in constant time, the average runtime is  $nc$ , where  $c$  is a constant (independent of  $\mathcal{L}$ ) and  $n - 1$  is the number of average recursion steps. Since in consecutive calls to the algorithm the variable  $i$  cycles through 0 to  $2^d - 1$  we conclude that the recursion is on average executed in every  $2^d$ th step, allowing for a geometric series argument.  $\square$

We see in Algorithm 4.3.10 the usefulness of the consecutive index. Because we are using this index instead of the TM-index, computing the index of the successor/predecessor only requires adding/subtracting 1 to the given index. On the other hand, computing the TM-index of a successor/predecessor would involve more subtle computations.

## 4.4. High-level AMR algorithms

To develop the complete AMR functionality required by numerical applications, we aim at a forest of quad-/octrees in the spirit of [34, 70]. Key high-level algorithms are (see also Section 2.5.2):

- **New.** Given an input mesh of conforming simplices, each considered a root simplex, generate an initial partitioned uniform refinement.
- **Adapt.** Adapt (refine and coarsen) a mesh according to a given criterion.
- **Partition.** Repartition a mesh among all processes such that the load is balanced, possibly according to weights.
- **Ghost.** For each process, assemble the layer of directly neighboring elements owned by other processes.
- **Balance.** Establish a 2:1 size condition between neighbors in a given refined mesh. Thus, the levels of any two neighboring simplices must differ by at most 1.
- **Iterate.** Iterate through the local mesh, executing a callback function on each element and on all inter-element interfaces.

Since partitioning via SFC only uses the SFC index as information, we refer to already existing descriptions of **Partition** for hexahedral or simplicial SFCs; see [34, 116] and Section 3.5. **Ghost** and **Balance** are sophisticated parallel algorithms and require additional theoretical work. We describe them in Chapters 7 and 8. For **Iterate** see our remarks in Section 9.

Here, we briefly describe **New** and **Adapt**. In the forest-of-trees approach we model an adaptive mesh by a coarse mesh of level 0  $d$ -simplices, the **trees**. Such a coarse mesh could be specified manually for simple geometries, or obtained from executing a mesh generator. Each level 0 simplex is identified with the root simplex  $T_d^0$  and then refined adaptively to produce the fine and potentially nonconforming mesh of  $d$ -simplices. These simplices are partitioned among all processes; thus each process holds a range of trees, of which the first and last may be incomplete: Their leaves are divided between multiple processes.

An entity  $\mathcal{F}$  of the structure **forest** consists of the following entries

- $\mathcal{F}.C$  — the coarse mesh;
- $\mathcal{F}.K$  — the process-local trees;
- $\mathcal{F}.E_k$  — for each local tree  $k$  the list of process-local simplices in tetrahedral Morton order.

We acknowledge that **New** and **Adapt** are essentially communication-free, but still serve well to exercise some of the fundamental algorithms described earlier.

#### 4.4.1. New

The **New** algorithm creates a partitioned uniform level  $\ell$  refined forest from a given coarse mesh. To achieve this, we first compute the first and last  $d$ -simplices belonging to the current process  $p$ . From this range we can calculate which trees belong to  $p$  and for each of these trees, the consecutive index of the first and last  $d$ -simplices on this tree. We then create the first simplex in a tree by a call to **T** (Algorithm 4.3.8). In contrast to the **New** algorithm in [34] we create the remaining simplices by calls to **Successor** instead of **T** to avoid the  $\mathcal{O}(\ell)$  runtime of **T** in the case of simplices. Our numerical tests, displayed in Figure 4.11, show that the runtime of **New** is in fact linear in the number of elements and does not depend on the level  $\ell$ . Within the algorithm,  $K$  denotes the number of trees in the coarse mesh and  $P$  the number of processes.

---

#### Algorithm 4.4.1: New(Coarse Mesh $C$ , int $\ell$ )

---

```

Result: A partitioned uniform level  $\ell$  forest with  $C$  as coarse mesh.
1  $n \leftarrow 2^{d\ell}$ ,  $N \leftarrow nK$  /*  $d$ -simplices per tree and global number of  $d$ -simplices */
2  $g_{\text{first}} \leftarrow \lfloor Np/P \rfloor$ ,  $g_{\text{last}} \leftarrow \lfloor N(p+1)/P \rfloor - 1$  /* Global numbers of first and last..
   */
3  $k_{\text{first}} \leftarrow \lfloor g_{\text{first}}/n \rfloor$ ,  $k_{\text{last}} \leftarrow \lfloor g_{\text{last}}/n \rfloor$  /* ..local simplex and local tree range */
4 for  $t \in \{k_{\text{first}}, \dots, k_{\text{last}}\}$  do
5    $e_{\text{first}} \leftarrow (t = k_{\text{first}}) ? g_{\text{first}} - nt : 0$ 
6    $e_{\text{last}} \leftarrow (t = k_{\text{last}}) ? g_{\text{last}} - nt : n - 1$ 
7    $T \leftarrow \mathbf{T}(e_{\text{first}}, \ell)$  /* Call Algorithm 4.3.8 */
8    $\mathcal{E}_k \leftarrow \{T\}$ 
9   for  $e \in \{e_{\text{first}}, \dots, e_{\text{last}} - 1\}$  do
10     $T \leftarrow \mathbf{Successor}(T)$ 
11     $\mathcal{E}_k \leftarrow \mathcal{E}_k \cup \{T\}$ 

```

---

After **New** returns, the process local number of elements is known, and per-element data can be allocated linearly in an array of structures, or a structure of arrays, depending on the specifics of the application.

We point out, that the operations that require specific knowledge of simplices are outsourced to the low-level algorithms **T** and **Successor**. If we replace, for example, the implementation of these by the appropriate versions for quadrilaterals/hexahedra with the Morton index, we obtain a partitioned uniform mesh with these elements. Even further, we can model hybrid meshes, when we store the low-level functions as part of the tree information, i.e. each tree has its own element type.

Thus, this description of **New** fits into our general approach of separating high- and low-level algorithms in order to handle different element types, which we describe in Section 2.5.3

### 4.4.2. Adapt

The **Adapt** algorithm modifies an existing forest by refining and coarsening the  $d$ -simplices of a given forest according to a callback function. It does this by traversing the  $d$ -simplices of each tree in tetrahedral Morton order and passing them to the callback function. If the current  $d$ -simplex and its  $2^d - 1$  successors form a family (all having the same parent), then the whole family is passed to the callback. This callback function accepts either one or  $2^d$   $d$ -simplices as input plus the index of the current tree. In both cases, a return value greater than zero means that the first input  $d$ -simplex should be refined, and thus its  $2^d$  children are added in tetrahedral Morton order to the new forest. Additionally, if the input consists of  $2^d$  simplices, they form a family, and a return value smaller than zero means that this family should be coarsened, thus replaced by their parent. If the callback function returns zero, the first given  $d$ -simplex remains unchanged and is added to the new forest, and **Adapt** continues with the next  $d$ -simplex in the current tree. The **Adapt** algorithm creates a new forest from the given one and can handle recursive refinement/coarsening. For the recursive part we make use of the following reasonable assumptions:

- A  $d$ -simplex that was created in a refine step will not be coarsened during the same adapt call.
- A  $d$ -simplex that was created in a coarsening step will not be refined during the same adapt call.

From these assumptions we conclude that for recursive refinement we only have to consider those  $d$ -simplices that were created in a previous refinement step and that we only have to care about recursive coarsening directly after we processed a  $d$ -simplex that was not refined and could be the last  $d$ -simplex in a family. If refinement and coarsening are not done recursively, the runtime of **Adapt** is linear in the number of  $d$ -simplices of the given forest.

An application will generally project or otherwise transform data from the previous to the adapted mesh. This can be done within the adaptation callback, which is known to proceed linearly through the local elements, or after **Adapt** returns if a copy of the old mesh has been retained. In the latter case, one would allocate element data for the adapted mesh and then iterate over the old and the new data simultaneously, performing the projection in the order of the SFC. Once this is done, the old data and the previous mesh are deallocated [28].

## 4.5. Performance evaluation

Given the design of the algorithms discussed in this chapter, we expect runtimes that are precisely proportional to the number of elements and independent of the level of refinement. To verify this, we present scaling and runtime tests<sup>1</sup> for **New**

<sup>1</sup>Version v0.1 is available at <https://bitbucket.org/cburstedde/t8code.git>

and **Adapt** on the JUQUEEN supercomputer at the Forschungszentrum Jülich [73], an IBM BlueGene/Q system with 28,672 nodes consisting of 16 IBM PowerPC A2 @ 1.6 GHz and 16 GB Ram per node. We also present one runtime study on the full MIRA system at the Argonne Leadership Computing Facility, which has the same architecture as JUQUEEN and 49,152 nodes. The biggest occurring number of mesh elements is around  $850e9$  tetrahedra with 13 million elements per process.

The first two tests are a strong scaling (up to 131k processes) and a runtime study of **New** in 3D, shown in Figure 4.11. For both tests we use a coarse mesh of 512 tetrahedra. We time the **New** algorithm with input level 8 (resp. level 10 for higher numbers of processes). We execute the runtime study to examine whether **New** has the proposed level-independent linear runtime in the number of generated tetrahedra, which can be read from the results presented in the Table in Figure 4.11.

The last test is **Adapt** with a recursive nonuniform refinement pattern. The starting point for all runs is a mesh obtained by uniformly refining a coarse mesh of 512 tetrahedra to a given initial level  $k$ . This mesh is then refined recursively using a single **Adapt** call, where only the tetrahedra of types 0 and 3 whose level does not exceed the fine level  $k + 5$  are refined recursively. The resulting mesh on each tetrahedron resembles a fractal pattern similar to the Sierpinski tetrahedron. We perform several strong and weak scaling runs on JUQUEEN starting with 128 processes and scaling up to 131,072. The setting is 16 processes per compute node. We finally do another strong scaling run on the full system of the MIRA supercomputer at the Argonne Leadership Computing Facility with 786,432 processes and again 16 processes per compute node. Figure 4.12 shows our runtime results.

## 4.6. Conclusion

We present a new encoding for adaptive nonconforming triangular and tetrahedral mesh refinement based on Bey’s red-refinement rule. We identify six different types of tetrahedra (and two types of triangles) and prescribe an ordering of the children for each of these types that differs from Bey’s original order. By introducing an embedding of the mesh elements into a Cartesian coordinate structure, we define a tetrahedral Morton index that can be computed using bitwise interleaving similar to the Morton index for cubes. This tetrahedral Morton index shares some properties with the well-known cubical one and allows for a memory-efficient random access storage of the mesh elements.

Exploiting the Cartesian coordinate structure, we develop several constant-time algorithms on simplices. These include computing the parent, the children, and the face-neighbors of a given mesh element, as well as computing the next and previous elements according to the SFC.

In view of providing a complete suite of parallel dynamic AMR capabilities, the constructions and algorithms described in this chapter are just the beginning. A repartitioning algorithm following our SFC, for example, is easy to imagine, but

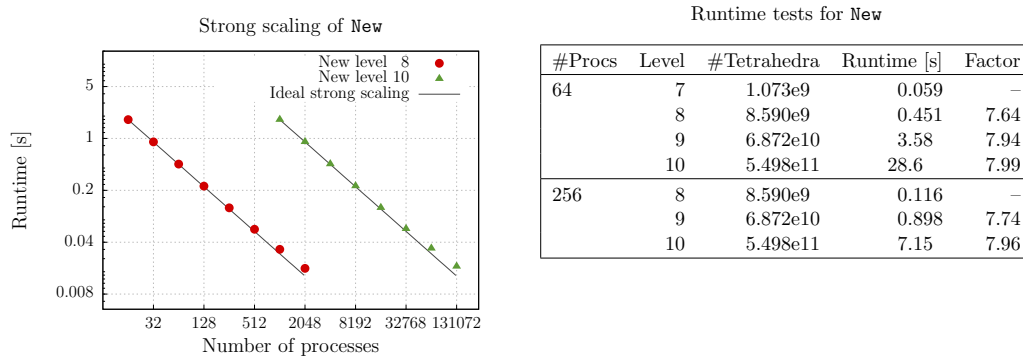


Figure 4.11.: Runtime tests for **New** on JUQUEEN. Left: Two strong scaling studies. A new uniform level 8 (circles) and level 10 (triangles) refinement of a coarse mesh of 512 root tetrahedra, carried out with 16 up to 2,048 processes and 1,024 up to 131,072 processes with 16 processes per compute node. Right: The data shows that the runtime of **New** is linear in the number of generated elements and does not additionally depend on the level. The uniform refinement is created from a coarse mesh of 512 root tetrahedra. For the first computation on 64 processes we use 1 process per compute node and for the computation on 256 processes we use 2 processes per node.




Figure 4.12.: Strong scaling for **Adapt** with a fractal refinement pattern. Starting from an initial level  $k$  on a coarse mesh of 512 tetrahedra we refine recursively to a maximal final level  $k + 5$ . The refinement callback is such that only subtetrahedra of types 0 and 3 are refined. Left: Strong and weak scaling on JUQUEEN with up to 131,072 processes and strong scaling on MIRA with up to 786,432 processes. On both systems we use 16 processes per compute node. The level 12 mesh consists of 858,588,635,136 tetrahedra. Right: An initial level 0 and final level 3 refinement according to the fractal pattern. The subtetrahedra of levels 1 and 2 are transparent.

challenging to implement if the tree connectivity is to be partitioned dynamically, and if global shared metadata shall be reduced from being proportional to the number of ranks to the number of compute nodes, see Chapter 6. The present chapter provides atomic building blocks that can be used in high-level algorithms for 2:1 balancing [69] and the computation of ghost elements and generalized topology iteration [70]. We address these algorithms in Chapters 7 and 8. The choices presented in this chapter are sustainable for maintaining extreme scalability in the long term.

## 5. Connected Components of the TM-SFC

This chapter is based on the preprint [32]. Since the preprint also contains contributions of Burstedde and Isaac, we only present those parts of it that are work of the author of this thesis. In particular, these are the results about the number of face-connected components for the TM-SFC.

When we store a mesh in parallel, its elements are distributed among the processes along the SFC. This means that each process is assigned a range of elements that is contiguous with regards to their SFC index. See also Section 3.5.

In many numerical applications, the processes communicate with each other across the boundaries of their partitions. See for example the construction of a ghost layer in Chapter 7. The total volume of parallel communication is thus proportional to the number of processes that share boundaries of a process' domain and proportional to the number of elements at this boundary. Hence, in order to minimize the communication, the surface-to-volume ratio of these partitions should be small. A good indicator for this is the number of face-connected components of a process' domain, and thus the number of face-connected components of a segment of the SFC.

It is a known fact that the cubical Morton SFC can produce disconnected segments. However, the number of face-connected components was shown to be at most two [9, 32]. Similarly, the TM-SFC can produce disconnected domains (see for example Figure 5.1). In this chapter, we prove the following bounds for the count of face-connected components.

**Theorem 5.1.** *A contiguous segment of a tetrahedral Morton curve through a uniform or adaptive tree of maximum refinement level  $L \geq 2$  produces at most  $2(L - 1)$  face-connected subdomains in 2D and at most  $2L + 1$  in 3D. For  $L = 1$  there are at most two face-connected subdomains.*

We complete our study with numerical results to illustrate the distribution of continuous vs. discontinuous segments. This supports the conjecture that the tetrahedral Morton curve is no worse in practice than the original cubical construction.

### 5.1. Proof of Theorem 5.1

We examine the number of face-connected components of a segment of the tetrahedral Morton SFC. As we show in Figure 5.1, there exist cases where the number



of face-connected components in a uniform 2D level  $L$  refinement can be as high as  $2(L - 1)$ . We show that this is in fact a sharp upper bound. We also show that in three dimensions the number of face-connected components does not exceed  $2L + 1$ . There exists an example with  $2L$  face-connected components and we conjecture that  $2L$  is in fact the sharp estimate. The proof of these bounds is fairly analogous to the results for cubes from [32] and relies on a divide-and-conquer approach by splitting the segment into subsegments of which we know the number of face-connected components.

**Remark 5.2.** In this chapter we sometimes refer to a TM-SFC on a type 1 root triangle. By this we mean the SFC that would result from the construction in Section 4.2 if we took the triangle  $2^{\mathcal{L}}S_1$  as root triangle; see also Figure 4.2. We use it since the geometry of this curve is the same as if we consider the original (type 0) TM-SFC and restrict it to the level 1 subtriangle of type 1.

**Remark 5.3.** For the cubical Morton curve, the proof that there are at most two face-connected components uses a symmetry property of the Morton index (see [32, equation (4)]). Given a  $d$ -dimensional cube  $q$  of level  $L$  with Morton index  $Q$ , we can take the bitwise negation of  $Q$

$$R(Q) = 2^{dL} - 1 - Q \tag{5.1}$$

to obtain the index of a new cube  $q'$ . This  $q'$  is the cube that results from traversing the SFC  $Q$  steps in reverse.

We do not have such a strong symmetry property for simplices. However, in 2D it holds that reversing the TM curve in a uniform refinement of a type 0 root triangle results in the forward TM curve for a type 1 root triangle.

Using this Remark, we show a weaker form of Corollary 8 from [32].

**Lemma 5.4.** *The following two properties hold for the TM-index in 2D, where we consider a uniform level  $L$  refinement of an initial type 0 triangle  $T$ .*

- *Each type 1 subsimplex is face-connected to a type 0 subsimplex with a greater TM-index.*
- *Each type 0 subsimplex that is also a descendant of the level 1, type 1 subtriangle  $T_3$  is face-connected to a type 1 subsimplex with a greater TM-index.*

*Proof.* The respective face-neighbor is the top face-neighbor for the type 1 subsimplex and the face-neighbor along the diagonal face for the type 0 subsimplex; see Figure 5.2. For type 0 we additionally require that the subsimplex is a descendant of  $T_3$ , since this ensures that the face-neighbor along the diagonal face is inside the root triangle. Despite this detail, the proofs for both items are identical, and we only present one for the first.

Let  $S$  denote an arbitrary type 1 subsimplex of level  $L$  and let  $S'$  be its neighbor across the top face. If  $S$  and  $S'$  share the same parent  $P$  then there are two cases,




Figure 5.1.: Left: A segment of the 2D SFC on a level 4 refinement of  $T_2^0$  with six face-connected components (shaded pink). The number of face-connected components in 2D can be as high as  $2(L - 1)$ ; this estimate is sharp. Right: a 3D level 2 refinement of  $T_3^0$  with four ( $= 2L$ ) face-connected components. We prove that an upper bound on the number of face-connected components is  $2L + 1$  and conjecture that  $2L$  is sharp.

which we also see in Figure 4.7: Either  $\text{type}(P) = 0$ , then the local index of  $S$  is 2 and that of  $S'$  is 3, or  $\text{type}(P) = 1$ , in which case the local index of  $S$  is 0 and that of  $S'$  is 1. Thus, in both cases the TM-index of  $S$  must be smaller than that of  $S'$ . We suppose now that  $S$  and  $S'$  have different parents, which implies  $L \geq 2$ , and denote these different level  $L - 1$  subsimplices by  $P$  and  $P'$ . The only possible combination is that  $\text{type}(P) = 1$  and  $\text{type}(P') = 0$ , and that  $P$  and  $P'$  are neighbors along  $P$ 's top face. Therefore, by an induction argument,  $m(P) < m(P')$ , and since the TM-index preserves the local order under refinement, each child of  $P$  has a smaller TM-index than each child of  $P'$ ; see Theorem 4.18. In particular we find  $m(S) < m(S')$ .  $\square$

We can now prove a first connectivity result for the 2D case.

**Lemma 5.5.** *Consider a triangle  $T$  that is uniformly refined to level  $L$ . If  $T$  has type 0, then a contiguous segment of the SFC ending in the last level  $L$  subsimplex has just one face-connected component. If  $T$  has type 1, then this holds for segments starting in the first level  $L$  subsimplex.*

*Proof.* We present the proof for  $\text{type}(T) = 0$ , since we can then use the symmetry of the 2D curve (Remark 5.3) to obtain the result for the case  $\text{type}(T) = 1$ . We proceed by induction over  $L$ .

For  $L = 0$  there is only one possible segment and it is face-connected. For  $L = 1$  we obtain the result by investigating all 10 cases. For  $L > 1$ , let  $j \in \{0, 1, 2, 3\}$  be the local index of the level 1 subtree  $T'$  of  $T$  in which the first level  $L$  subsimplex of the segment lies. If  $j \in \{0, 1, 3\}$ , then the type of  $T'$  is 0 and the statement follows by induction. Thus, let  $j = 2$ , i.e., the segment starts in the type 1 subtree of  $T$ . The part of the segment that is not inside  $T'$  is the full last subtree of  $T$  (local index 3) and thus it is face-connected in itself. With Lemma 5.4 we conclude that




Figure 5.2.: Illustration of Lemma 5.4. In 2D, choose any subsimplex  $S_*$ . If its neighbor along the top face  $S'_*$  is inside the root triangle, then  $m(S_*) < m(S'_*)$ . This condition is always fulfilled by any type 1 triangle and by type 0 triangles that are descendants of the middle level 1 subtriangle.

each subsimplex in the subsegment in  $T'$  is face-connected to a simplex with greater TM-index. Iterating this process, we conclude that each of these subsimplices is face-connected to a subsimplex of the full last subtree of  $T$ . Thus, the whole segment is face-connected.  $\square$

For all other segments beginning with the first or ending in the last level  $L$  subsimplex, and notably for all of those segments in 3D, we obtain an upper bound of  $L + 1$  face-connected components, which we show in the next two lemmas.

**Lemma 5.6.** *Let a segment of the TM-SFC for a uniform level  $L$  refined  $d$ -simplex consist of several full level 1 subsimplices plus one single level  $L$  simplex either at the end or at the beginning, then this segment has at most two face-connected components.*

*Proof.* Similarly to the proof of Proposition 9 in [32], we can show this claim by enumerating all possible cases. There is no induction necessary.  $\square$

**Lemma 5.7.** *If a  $d$ -simplex is uniformly refined to level  $L$ , then any segment of the TM-SFC ending in the last subsimplex or starting in the first has at most  $L + 1$  face-connected components.*

*Proof.* Consider the case that the segment starts in the first simplex. For  $L = 0$  there is only one possible segment consisting of the unique level 0 subsimplex and it is thus face-connected. Let now  $L > 0$ . Since the segment begins at the very first level  $L$  subsimplex, we can separate it into two parts. The first part at the beginning consists of 0 to  $2^d - 1$  full level 1 subtrees, and the second part is one possibly incomplete level 1 subtree.

By the induction assumption, the second part has at most  $L$  face-connected components. From Lemma 5.6 we obtain that the first part together with the first level

$L$  subsimplex of the second part has at most two face-connected components. Since this first level  $L$  subsimplex is contained in one of the components of the second part, we obtain

$$L + 2 - 1 = L + 1 \tag{5.2}$$

components in total.

If the segment ends in the last simplex, the order of parts is reversed. The first part of the segment is the part in the level 1 subtree where the segment starts, and the second part consists of the remaining full level 1 subtrees. We obtain the bound on the number of face-connected components using the same inductive reasoning as above.  $\square$

We have so far argued the connectivity of specific kinds of SFC segments. This suffices to proceed to arbitrary segments of the tetrahedral Morton SFC.

**Proposition 5.8.** *Any contiguous segment of the TM-SFC of a uniform level  $L \geq 2$  refinement of a type 0 simplex has at most  $2(L - 1)$  face-connected components in 2D and  $2L + 1$  face-connected components in 3D. For  $L = 1$ , there are at most two face-connected components, and one for  $L = 0$  (this applies to both 2D and 3D).*

*Proof.* Again, the cases  $L = 0$  and  $L = 1$  follow by inspecting all cases. Thus, let  $L \geq 2$ . We first show that for  $d \leq 3$  the number of face-connected components is bounded by  $2L + 1$ : If a given segment is contained in a level 1 subtree, we are done by induction. Otherwise we can divide the segment into three (possibly empty) pieces: First, the segment in one incomplete level 1 subtree ending at its last level  $L$  subsimplex, then one contiguous segment of full level 1 subtrees and finally a segment in one (possibly incomplete) level 1 subtree that starts at its first level  $L$  subsimplex. Lemma 5.7 implies that the first and the last piece have at most  $L$  face-connected components each. By Lemma 5.6, the second piece has one or two face-connected components, and if the number is two, then it is face-connected to the first or to the third piece. Thus, it adds only one face-connected component to the total number, and we obtain at most

$$L + 1 + L = 2L + 1 \tag{5.3}$$

face-connected components.

Let us now specialize to 2D. We conclude from Lemma 5.5 that the first subsegment only adds more than one face-connected component if it is contained in the only level 1 subtree of type 1 (local index 2). Similarly, the third subsegment only adds more than one face-connected component if it is contained in a level 1 subtree of type 0. In particular, if both subsegments add more than one face-connected component, the third subsegment is contained in the last level 1 subtree (local index 3). Thus, the second subsegment is empty in this case.

If both of these subsegments have less than  $L$  face-connected components, there is nothing left to show since the overall number of components is then less than or equal to  $2(L - 1)$ . So suppose that one of the subsegments has  $L$  face-connected




Figure 5.3.: An illustration of the 2D case in the proof of Proposition 5.8 for  $L = 4$ . The bottom segment has the maximal number of  $L$  face-connected components. Since its first and second triangle (on the left, shaded in pink) are face-connected with the top segment, the possible number of face-connected components is reduced by two. If the second segment has  $L$  face-connected components as well, then its last two triangles (on the right) are face-connected with the bottom segment. Thus, the number of face-connected components is less than or equal to  $2L - 2$ .

components and the other one has at least  $L - 1$ . We depict this situation in Figure 5.3. We observe that the first and second level  $L$  simplex in this first segment are face-connected to the first and second level  $L$  simplex in the second segment. If, however, the second subsegment has  $L$  face-connected components then its last two level  $L$  simplices are face-connected to the last two level  $L$  simplices of the first subsegment.

We thus can subtract two face-connected components from the total count, which leads to at most

$$L + L - 2 = 2(L - 1) \tag{5.4}$$

face-connected components in total. □

We briefly discuss whether we can sharpen these bounds. In 2D, this is not possible by counterexample; see Figure 5.1. In 3D, we construct a segment with  $2L$  face-connected components using the consecutive SFC-indices 22–25 of a uniform level 2 refinement of a type 0 tetrahedron. We believe that the case that the first and the last piece described in the proof of Proposition 5.8 have  $L$  face-connected components each and that additionally the middle piece adds one component does not occur.

**Conjecture 5.9.** *In 3D, the number of face-connected components is bounded by  $2L$ . This estimate is sharp.*

### 5.1.1. From uniform to adaptive meshes

We close this section with the extension of the proof from uniform to adaptive meshes, which is the remaining step to establish Theorem 5.1,

We have completed the necessary proofs for a uniform space division for triangular and tetrahedral refinement (see 5.1). As we state in this section, an adaptive space division does not require any more effort (see also [32] and [9, page 176]).

*Proof of Theorem 5.1.* Any adaptive tree of simplices with level  $l \leq L$  can be refined into level  $L$  simplices exclusively. This operation does not change the connectivity between boundaries of the designated subdomain. In particular, the number of face-connected subdomains remains unchanged and the proof reduces to applying Proposition 5.8 above.  $\square$

### 5.1.2. From one tree to a forest

If we consider a forest of octrees as in Section 3.4 or in [10, 34, 134], a contiguous segment of the TM curve may traverse more than one tree. In this case, the segment necessarily contains the last subsimplex of any predecessor tree, as well as the first subsimplex of any successor tree in the segment. For the simplicial case, we may use Lemmas 5.5 and 5.7 to use the bounds  $L + 1$  (2D) and  $2L + 1$  (3D) for the respective parts of the segment, not counting the transitions between full trees.

## 5.2. Enumeration of face-connected segments

We would like to examine not only how many pieces an SFC segment can have, but also how frequently segments of different numbers of pieces occur. To this end, we supply numerical studies for the TM-SFCs and compare the results with the cubical Morton curve.

We enumerate all possible TM-SFC segments for a given uniform refinement level and compute the number of their face-connected components. We also compute the relative counts of face-connected and non-connected segments for the cubical Morton curves. We achieve this by performing a depth-first search on the connectivity graph of the submesh generated by the segment<sup>1</sup>.

In an application, all possible lengths of SFC segments can occur. On the one hand, we could have a forest consisting of a single tree. If the number of participating processes is of the same magnitude than the number of elements in that tree, then very small lengths of segments occur, possibly even segments consisting only of a single element. On the other hand, consider a setting where we have many trees, possibly as many or more trees than processes. In this case, the lengths of SFC segments within a single tree can be arbitrarily large, reaching up to the maximum of the full tree. See for example our discussion in Section 6.4.

---

<sup>1</sup>[https://github.com/holke/sfc\\_conncomp](https://github.com/holke/sfc_conncomp)

	Level 5				Level 8	
	Quads	Cubes	Triangles	Tets	Quads	Triangles
Face-connected	71.6%	60.0%	63.9%	61.0%	71.4%	63.7%
Non-connected	28.4%	40.0%	36.1%	39.0%	28.6%	36.3%

Table 5.1.: The relative counts of face-connected and non-connected segments across all possible SFC segments of a uniform level 5 and level 8 (2D only) refinement.

We now compute the fraction of face-connected segments of any length among all possible segments. More precisely, we compute for each possible count of face-connected components the chance that any randomly chosen SFC segment (with a random length) has exactly this number of face-connected components.

For a uniform level 5 refined tetrahedron we obtain that 61% of all SFC segments are face-connected and only 7% have four or more face-connected components. For a uniform level 8 refined triangle, about 64% of the segments are face-connected with 2% of the segments having four or more components. For cubes and quadrilaterals, the respective ratios of face-connected segments are 60% and 71% (here we know that the disconnected segments have exactly two components). Thus, comparing cubical and TM curve, we see that in 2D more segments of the quadrilateral Morton curve are face-connected, and in 3D more segments of the TM curve are face-connected.

We collect these results in Figure 5.4 and Table 5.1.

### 5.3. Conclusion

We show that the bound for the TM-SFC is of order  $L$  and thus growing with the level of refinement. Yet, we can demonstrate numerically that the fraction of face-connected to non-connected segments is close to the cubical case. In practice, we may expect both approaches to behave similarly.

Our result would appear relevant to make informed choices about the type of space-filling curve to use, for example in writing a new element-based parallel code for the numerical solution of partial differential equations, or any other code that benefits from a recursive subdivision of space. Our theory and experiments support the existing numerical evidence that a fragmentation of the parallel partition is not observed.




Figure 5.4.: The relative count of SFC segments (left  $y$ -axis) by number of face-connected components and the average length (right  $y$ -axis) of these segments. Left: the distribution for a uniform level 8 refined triangle. We observe that almost 98% of all SFC segments have three face-connected components or less. 63.7% are face-connected, 29.7% have two face-connected components and 4.4% have three face-connected components. Right: the distribution for a uniform level 5 refined tetrahedron. Here, more than 93% of the segments have three face-connected components or less. 61.0% have exactly one face-connected component, 22.1% two face-connected components and 10.7% three face-connected components. The highest number of segments occurring are  $14 = 2(8-1)$  in 2D, and  $10 = 2 \cdot 5$  in 3D. This is in agreement with Proposition 5.8 (2D) and Conjecture 5.9 (3D).



## 6. Coarse Mesh Partitioning

This chapter is based on the paper [31]. We edited it slightly in order to fit into the general notations of this thesis, without changing its mathematical content. Copyright © by SIAM. Unauthorized reproduction of this chapter is prohibited.

As we discuss in Chapters 2 and 3, a technique to model complex domain shapes is to patch multiple trees together in an unstructured coarse mesh, giving rise to a forest of elements. The number of trees per process is limited by the available memory to roughly  $1e5$  to  $1e6$  [34]. In industrial and medical meshing however, numbers in the range of one billion or more trees are not uncommon [52, 68, 119]. In order to support such cases, we need to partition the coarse mesh among the parallel processes.

Two main approaches for partitioning a forest of elements have been discussed [157], namely (a) assigning each tree and thus all of its elements to one owner process [24, 71, 126], or (b) allowing a tree to contain elements belonging to multiple processes [10, 34]. The first approach offers a simpler logic but may not provide acceptable load balance when the number of elements differs vastly between trees. The second allows for perfect partitioning of elements by number (the local numbers of elements between processes differ by at most one) but presents the issue of trees that are shared between multiple processes.

We choose paradigm (b) for speed and scalability, challenging ourselves to solve an  $n$ -to- $m$  communication problem for every coarse mesh element. Thus the objective of this chapter is to develop how to do this without handshaking (i.e., without having to determine separately which process receives from which) and with a minimal number of senders, receivers, and messages. Our main contribution is to avoid identifying a single owner process for each tree and instead treat all its sharer processes as algorithmically active, under the premise that they produce a disjoint union of the information necessary to be transferred. In particular, each process shall store the relevant tree meta data to be readily available, eliminating the need to transfer this data from a single owner process.

In this chapter, we also integrate the parallel transfer of ghost trees. The reason for this is that each process will eventually collect ghost elements, i.e., remote elements adjacent to its own. We discuss the ghost algorithm on elements in Chapter 7. Ghost elements of any process may be part of trees that are not in its local set. To disconnect the ghost element transfer from identifying and transferring ghost trees, we perform the latter as part of the coarse mesh partitioning, presently across tree faces. We study in detail what information we must maintain to reference neighbor trees of ghost trees (that may themselves be either local, ghost, or neither) and

propose an algorithm with minimal communication effort.

We have implemented the coarse mesh partitioning for triangles and tetrahedra using the TM-SFC designed in Chapter 4, and for quadrilaterals and hexahedra exploiting the logic from [34]. To demonstrate that our algorithms are safe to use, we verify that (a) small numbers of trees require run times on the order of milliseconds and thus present no noticeable overhead compared to a serial coarse mesh, and (b) the coarse mesh partitioning adds only a fraction of run time compared to the partitioning of the forest elements, even for extraordinarily large numbers of trees. We show a practical example of 3D dynamic AMR on  $8e3$  processes using  $383e6$  trees and up to  $25e9$  elements. To investigate the ultimate limit of our algorithms, we partition coarse meshes of up to  $371e9$  trees on JUQUEEN using  $917e3$  processes, obtaining a total run time of about 1.2s and a rate of  $340e3$  trees per second per process. On  $131e3$  processes with half as many ranks per compute node we obtain rates as high as  $750e3$  trees per second per process.

We may summarize our results by saying that partitioning the trees can be made even less costly than partitioning the elements and often executes so fast that it does not make a difference at all. This allows a forest code that partitions both trees and elements dynamically to treat the whole continuum of forest mesh scenarios, from one tree with nearly trillions of elements on the one extreme to billions of trees that are not refined at all on the other, with comparable efficiency.

## 6.1. Tree-based AMR

We repeat some of the concepts from Chapters 2 and 3.

A forest  $\mathcal{F}$  consist of a coarse mesh of trees and a fine mesh of elements that resides from refining the trees. We use SFCs to order the elements within each tree. SFCs map the  $d$ -dimensional elements of a refinement tree to an interval by assigning a unique integer index  $\mathcal{I}_{\mathcal{F}}(E)$  to each element  $E$ , see Lemma 3.21. Thus, we can order all elements of that refinement tree linearly in an array. As in [136], we do not store the internal (non-leaf) nodes of the tree.

The choice of SFC affects the ordering of these elements of the forest mesh and thus the parallel partition of elements. Possibilities include, but are not limited to, the Hilbert, Peano, or Morton curves for quadrilaterals and hexahedra [65, 101, 113, 149], as well as the Sierpiński curve for triangles [8, 129] and the tetrahedral Morton curve for triangles and tetrahedra from Chapter 4.

As of equation (3.22), a global order of elements is established first by tree and then by their index with respect to an SFC (see also [9]): We enumerate the  $K$  trees of the coarse mesh by  $0, \dots, K - 1$  and call the number  $k$  of a tree its *global index*. With the global index we naturally extend the SFC order of the leaves: Let a leaf element of the tree  $k$  have SFC index  $I$  (within that tree); then we define the combined index  $(k, I)$ . This index compares to a second index  $(k', J)$  as

$$(k, I) < (k', J) \quad :\Leftrightarrow \quad k < k' \text{ or } (k = k' \text{ and } I < J). \quad (6.1)$$

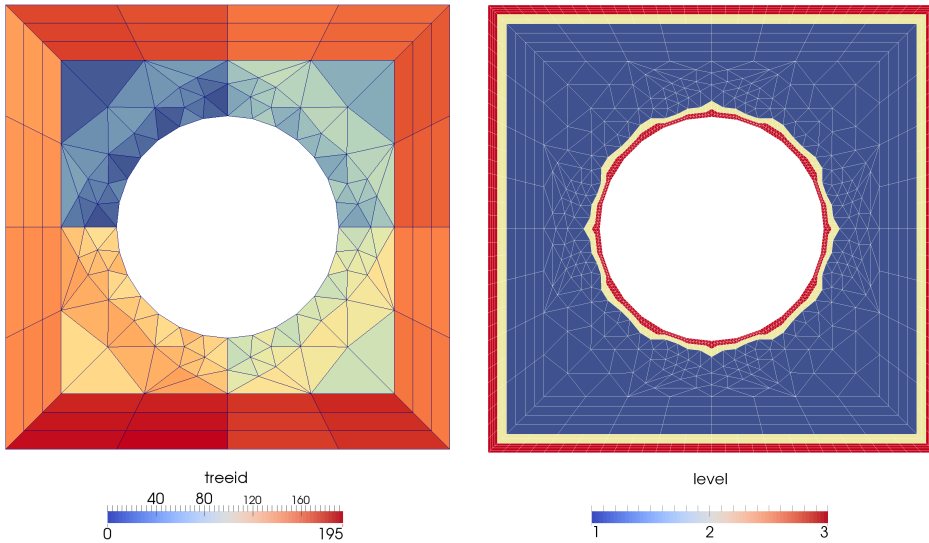


Figure 6.1.: A mesh consists of two structures, the coarse mesh (left) that represents the topology of the domain, and the forest mesh (right) that consists of the leaf elements of a refinement and is used for computation. In this example the domain is a unit square with a circular hole. The color coding in the coarse mesh displays each tree’s unique and consecutive identifier, while the color coding in the forest mesh represents the refinement level of each element. In this example we choose an initial global level 1 refinement and a refinement of up to level 3 along the domain boundary.

In practice we store the mesh elements local to a process in one contiguous array per locally nonempty tree in precisely this order.

The algorithms and techniques discussed in this chapter assume an SFC induced order among the elements, but they are not affected by the particular choice of SFC. In the `t8code` software used for the demonstrations in this thesis, we have so far implemented Morton SFCs for quadrilaterals and hexahedra via the `p4est` library [25] and the tetrahedral Morton SFC for tetrahedra and triangles. These curves compute the index  $m(E)$  of an element via bitwise interleaving the coordinates of its lower left vertex in a suitable reference tree; see also Figure 6.2 for an illustration of the curve on triangles. Other SFC schemes may be added to the `t8code` in a modular fashion, see Section 2.5.3.

### 6.1.1. The tree shapes

The trees of the coarse mesh can be of arbitrary shape as long as they are all of the same dimension and fit together along their faces. In particular, we identify the following tree shapes:




Figure 6.2.: We connect multiple trees to model complex geometries. Here, we show two trees  $k_0$  and  $k_1$  with an adaptive refinement. To enumerate the forest mesh, we establish an a priori order between the two trees and use an SFC within each tree. On the left-hand side of the figure the refinement tree and its linear storage are shown. When we partition the forest mesh to  $P$  processes (here,  $P = 3$ ), we cut the SFC in  $P$  equally sized parts and assign part  $i$  to process  $i$ .

Coarse mesh	Conforming mesh of tree roots
Tree	An element of the coarse mesh
Forest mesh	The adaptive mesh of elements (leaves of the trees)
Element/leaf	Each element of the forest mesh is the leaf of a tree

Table 6.1.: The basic definitions for the coarse mesh and the forest mesh and their elements. Throughout, we refer to the neighbor information of the trees as *connectivity*.

- Points in 0D.
- Lines in 1D.
- Quadrilaterals and triangles in 2D.
- Hexahedra and tetrahedra in 3D.
- Prisms and pyramids in 3D.

Coarse meshes consisting solely of prisms or pyramids are quite uncommon; these tree shapes are used primarily to transition between hexahedra and tetrahedra in hybrid meshes.

### 6.1.2. Encoding of face-neighbors

The connectivity information of a coarse mesh includes the neighbor relation between adjacent trees. Two trees are considered neighbors if they share at least one lower dimensional face (vertex, face, or edge). Since all of this connectivity information can be inferred from codimension-1 neighbors, we restrict ourselves to




Figure 6.3.: The vertex and face labels of the 2D (left) and 3D (right) tree shapes.

those, denoting them uniformly by face-neighbors. This choice does not lessen the generality of the partitioning algorithms to follow and avoids a significant jump in complexity of the element-neighbor code.

An application often requires a quick mechanism to access the face-neighbors of a given forest mesh element. If this neighbor element is a member of the same tree, the computation can be carried out via the SFC logic, which involves only a few bitwise operations for the hexahedral and tetrahedral Morton curves [29, 34, 101, 136]. If, however, the neighbor element belongs to a different tree, we need to identify this tree, given the parent tree of the original element and the tree face at which we look for the neighbor element. It is thus advantageous to store the face-neighbors of each tree in an array that is ordered by the tree's faces. To this end, we fix the enumeration of faces and vertices relative to each other as depicted in Figure 6.3. We discuss the computation of element face-neighbors in Chapter 7.

### 6.1.3. Orientation between face-neighbors

In addition to the global index of the neighbor tree across a face, we describe how the faces of the tree and its neighbor are rotated relative to each other. We allow all connectivities that can be embedded in a compact 2- or 3-manifold in such a way that each tree has positive volume. This includes the Moebius strip and Klein's bottle and other quite exotic meshes, e.g., a hexahedron whose one face connects to another in some rotation. We obtain two possible orientations of a line-to-line connection, three for a triangle-to-triangle connection, and four for a quadrilateral-to-quadrilateral connection.

We would like to encode the orientation of a face connection analogously to the way it is handled in `p4est`: At first, given a face  $f$ , its vertices are a subset of the vertices of the whole tree. If we order them accordingly and renumberate them consecutively starting from zero, we obtain a new number for each vertex that depends on the face  $f$ . We call it the *face corner number*. If now two faces  $f$  and  $f'$  meet, the corner 0 of the face with the smaller face number is identified with a



There are two traditional approaches to partition the coarse mesh. In the first approach [24, 152], the coarse mesh is partitioned, and the owner process of a tree will own all elements of that tree. In other words, it is not possible for two different processes to own elements of the same tree, which can lead to highly imbalanced forest meshes. Furthermore, if there are fewer trees than processes, there will be idle processes assigned zero elements. In particular, this approach prohibits the frequent special case of a single tree.

The second approach [157] is to first partition the forest mesh and then deduce the coarse mesh partition from that of the forest. If several processes have leaf elements from the same tree, then the tree is assigned to one of these processes, and whenever one of the other processes requests information about this tree, communication is invoked. This technique has the advantage that the forest mesh is load-balanced much better, but it introduces additional synchronization points in the program and can lead to critical bottlenecks if a lot of processes request information on the same tree.

We propose another approach, which is a variation of the second, that overcomes the communication issue. If several processes have leaf elements from the same tree, we duplicate this tree's connectivity data and store a local copy of it on each of the processes. Thus, there is no further need for communication, and each process has exactly the information it requires. Since the purpose of the coarse mesh is not to store data that changes during the simulation but to store connectivity data about the physical domain, the data on each tree is persistent and does not change during the simulation. Certainly, this concept poses an additional challenge in the (re)partitioning process, because we need to manage multiple copies of trees without producing redundant messages.

As an example, consider the situation in Figure 6.4. Here the 2D coarse mesh consists of two triangles 0 and 1, and the forest mesh is a uniform level 1 mesh consisting of 8 elements. Elements 0, 1, 2, 3 belong to tree 0 and elements 4, 5, 6, 7 to tree 1. If we load-balance the forest mesh to three processes with ranks 0, 1, and 2, then a possible forest mesh partition arising from an SFC could be

$$\begin{array}{ccc}
 \begin{array}{c} \text{rank} \quad \text{elements} \\ \hline 0 \quad 0, 1, 2 \\ 1 \quad 3, 4, 5 \\ 2 \quad 6, 7, \end{array} & \begin{array}{c} \text{leading to the coarse} \\ \text{(6.4) mesh partition} \end{array} & \begin{array}{c} \text{rank} \quad \text{trees} \\ \hline 0 \quad 0 \\ 1 \quad 0, 1 \\ 2 \quad 1. \end{array} \quad (6.5)
 \end{array}$$

Thus, each tree is stored on two processes.

### 6.2.1. Valid partitions

We allow arbitrary partitions for the forest, as long as they are induced by an SFC. This gives us some information on the type of coarse mesh partitions that we can expect.

**Definition 6.3.** In general, a **partition** of a coarse mesh of  $K$  trees  $\{0, \dots, K - 1\}$




Figure 6.4.: A coarse mesh of two trees and a uniform level 1 forest mesh. If the forest mesh is partitioned to three processes, each tree of the coarse mesh is requested by two processes. The numbers in the tree corners denote the position and orientation of the tree vertices, and the global tree ids are the numbers in the center of each tree. As an example SFC we take the triangular Morton curve from [29], but the situation of multiple trees per process can occur for any SFC.

to  $P$  processes  $\{0, \dots, P - 1\}$  is a map  $f$  that assigns each process a certain subset of the trees,

$$f: \{0, \dots, P - 1\} \longrightarrow \mathcal{P} \{0, \dots, K - 1\}, \quad (6.6)$$

and whose image covers the whole mesh:

$$\bigcup_{p=0}^{P-1} f(p) = \{0, \dots, K - 1\}. \quad (6.7)$$

Here,  $\mathcal{P}$  denotes the set of all subsets (power set). We call  $f(p)$  the *local trees* of process  $p$  and explicitly allow that  $f(p) \cap f(q)$  may be nonempty. If so, the trees in this intersection are *shared* between processes  $p$  and  $q$ .

The above definition includes all possibilities of attaching trees to processes. In particular, we allow the same tree to reside on multiple processes. For example, it is a partition in the above sense if all trees are partitioned to all processes, expressed by  $f(p) = \{0, \dots, K - 1\}$  for all  $p$ . Any other arbitrary mapping is possible, with no restriction on the number of trees per process, where empty processes with  $f(p) = \emptyset$  are included.

For our method, we will not consider every possible partition of a coarse mesh. Since we assume that a forest mesh partition comes from an SFC, we restrict ourselves to a subset of partitions. In particular, the SFC order of a forest imposes the restriction that the order of trees is linear among the processes. Thus if tree  $k$  is on process  $p$ , then every tree  $l > k$  must be partitioned to a process  $q \geq p$ .

**Definition 6.4.** Consider a partition  $f$  of a coarse mesh with  $K$  trees. We say that  $f$  is a **valid partition** if there exist a forest mesh with  $N$  leaves and a (possibly weighted) SFC partition of it that induces  $f$ . Thus, for each process  $p$  and each tree  $k$ , we have  $k \in f(p)$  if and only if there exists a leaf  $e$  of the tree  $k$  in the forest mesh that is partitioned to process  $p$ . Processes without any trees are possible; in this case  $f(p) = \emptyset$ .



We denote by  $k_p$  the local tree on  $p$  with the lowest global index and denote by  $K_p$  the local tree with the highest global index.

This, for example, excludes meshes where the processes are not mapped to the trees in ascending order. A simple example of a partition of two trees on two processes that is not valid is given by  $f(0) = \{1\}$  and  $f(1) = \{0\}$ .

The definition of valid partitions requires a forest mesh and a specific SFC-induced partition of it. Since this is not convenient for theoretical investigations, we deduce three properties that characterize valid partitions independently of a forest mesh and SFCs.

**Proposition 6.5.** *A partition  $f$  of a coarse mesh is valid if and only if it fulfills the following properties.*

(i) *The tree indices of a process's local trees are consecutive; thus*

$$f(p) = \{k_p, k_p + 1, \dots, K_p\} \text{ or } f(p) = \emptyset. \quad (6.8)$$

(ii) *A tree index of a process  $p$  may be smaller than a tree index on a process  $q$  only if  $p \leq q$ :*

$$p \leq q \Rightarrow K_p \leq k_q \quad (\text{if } f(p) \neq \emptyset \neq f(q)). \quad (6.9)$$

(iii) *The only trees that can be shared by process  $p$  with other processes are  $k_p$  and  $K_p$ :*

$$f(p) \cap f(q) \subseteq \{k_p, K_p\} \text{ for } p \neq q. \quad (6.10)$$

*Proof.* We show the only-if direction first. Let an arbitrary forest mesh with SFC partition be given such that  $f$  is induced by it. In the SFC order the leaves are sorted according to their SFC indices. If  $(i, I)$  denotes the leaf corresponding to the  $I$ -th leaf in the  $i$ -th tree and tree  $i$  has  $N_i$  leaves, then the complete forest mesh consists of the leaves

$$\{(0, 0), (0, 1), (0, N_0 - 1), (1, 0), \dots, (K - 1, N_{K-1} - 1)\}. \quad (6.11)$$

The partition of the forest mesh is such that each process  $p$  gets a consecutive range

$$\{(k_p, i_p), \dots, (K_p, i'_p)\} \quad (6.12)$$

of leaves, where  $(k_{p+1}, i_{p+1})$  is the successor of  $(K_p, i'_p)$  and the  $k_p$  and the  $K_p$  form increasing sequences with  $K_p \leq k_{p+1}$ . The coarse mesh partition is then given by

$$f(p) = \{k_p, k_p + 1, \dots, K_p\} \text{ for all } p, \quad (6.13)$$

which shows properties (i) and (ii). To show (iii) we assume that  $f(p)$  has at least three elements; thus  $f(p) = \{k_p, k_p + 1, \dots, K_p\}$ . However, this means that in the forest mesh partition each leaf of the trees  $\{k_p + 1, \dots, K_p - 1\}$  is partitioned to  $p$ .

Since the forest mesh partitions are disjoint, no other process can hold leaf elements from these trees, and thus they cannot be shared.

To show the if-direction, suppose the partition  $f$  fulfills (i), (ii), and (iii). We construct a forest mesh with a weighted SFC partition as follows. Each tree that is local to a single process is not refined and thus contributes a single leaf element to the forest mesh. If a tree is shared by  $m$  processes, then we refine it uniformly until we have more than  $m$  elements. It is now straightforward to choose the weights of the elements such that the corresponding SFC partition induces  $f$ .  $\square$

We directly conclude the following.

**Corollary 6.6.** *In a valid partition, each pair of processes can share at most one tree; thus*

$$|f(p) \cap f(q)| \leq 1 \quad (6.14)$$

for each  $p \neq q$ .

*Proof.* Supposing the contrary, with (6.10) we know that there would exist two processes  $p$  and  $q$  with  $p < q$  such that  $f(p) \cap f(q) = \{k_p, K_p\} = \{k_q, K_q\}$  and  $k_p \neq K_p$ . Thus  $K_p > k_p = k_q$ , which contradicts property (6.9).  $\square$

**Corollary 6.7.** *If in a valid partition  $f$  of a coarse mesh the tree  $k$  is shared between processes  $p$  and  $q$ , then for each  $p < r < q$ ,*

$$f(r) = \{k\} \quad \text{or} \quad f(r) = \emptyset. \quad (6.15)$$

*Proof.* We can directly deduce this from (6.8), (6.9), and Corollary 6.6.  $\square$

In order to properly deal with empty processes in our calculations, we define start and end tree indices for these as well.

**Definition 6.8.** Let  $p$  be an empty process in a valid partition  $f$ ; thus  $f(p) = \emptyset$ . Furthermore, let  $q < p$  be maximal such that  $f(q) \neq \emptyset$ . Then we define the start and end indices of  $p$  as

$$k_p := K_q + 1, \quad (6.16a)$$

$$K_p := K_q = k_p - 1. \quad (6.16b)$$

If no such  $q$  exists, then no rank lower than  $p$  has local trees, and we set  $k_p = 0$ ,  $K_p = -1$ . With these definitions, (6.8) and (6.9) are valid if any of the processes are empty.

From now on, all partitions in this chapter are assumed valid even if not stated explicitly.

## 6.2.2. Encoding a valid partition

A typical way to define a partition in a tree-based code is to store an array  $\mathbf{O}$  of tree offsets for each process, that is, the global index of the first tree local to each process. The range of local trees for process  $p$  can then be computed as  $\{\mathbf{O}[p], \dots, \mathbf{O}[p+1] - 1\}$ . However, for valid partitions in the coarse mesh setting, this information would not be sufficient because we would not know which trees are shared. We thus modify the offset array by adding a negative sign when the first tree of a process is shared.

**Definition 6.9.** Let  $f$  be a valid partition of a coarse mesh, with  $k_p$  being the index of  $p$ 's first local tree. Then we store this partition in an array  $\mathbf{O}$  of length  $P + 1$ , where for  $0 \leq p < P$ ,

$$\mathbf{O}[p] := \begin{cases} k_p & \text{if } k_p \text{ is not shared with the next smaller} \\ & \text{nonempty process or } f(p) = \emptyset, \\ -k_p - 1 & \text{if it is.} \end{cases} \quad (6.17)$$

Furthermore,  $\mathbf{O}[P]$  shall store the total number of trees.

Because of the definition of  $k_p$ , we know that  $\mathbf{O}[0] = 0$  for all valid partitions.

**Lemma 6.10.** Let  $f$  be a valid partition, and let  $\mathcal{O}$  be as in Definition 6.9. Then

$$k_p = \begin{cases} \mathcal{O}[p] & \text{if } \mathcal{O}[p] \geq 0, \\ |\mathcal{O}[p] + 1| & \text{if } \mathcal{O}[p] < 0, \end{cases} \quad (6.18)$$

and

$$K_p = |\mathcal{O}[p+1]| - 1. \quad (6.19)$$

*Proof.* The first statement follows since (6.17) and (6.18) are inverses of each other. For (6.19) we distinguish two cases. First, let  $f(p)$  be nonempty. If the last tree of  $p$  is not shared with  $p+1$ , then it is  $k_{p+1} - 1$  and  $\mathbf{O}[p+1] = k_{p+1}$ , and thus we have

$$K_p = k_{p+1} - 1 = |\mathbf{O}[p+1]| - 1. \quad (6.20)$$

If the last tree of  $p$  is shared with  $p+1$ , then it is  $k_{p+1}$ , the first local tree of  $p+1$ , and thus  $\mathbf{O}[p+1] = -k_{p+1} - 1$  and

$$K_p = k_{p+1} = |-k_{p+1}| = |\mathbf{O}[p+1]| - 1. \quad (6.21)$$

Now let  $f(p) = \emptyset$ . If  $k_{p+1}$  is not shared, then  $k_{p+1} = k_p = K_p + 1$  by Definition 6.8, and  $\mathbf{O}[p+1] = k_{p+1}$  by (6.17). Thus,

$$K_p = k_p - 1 = k_{p+1} - 1 = |\mathbf{O}[p+1]| - 1. \quad (6.22)$$

If  $k_{p+1}$  is shared, then again by Definition 6.8,  $k_{p+1} = k_p - 1 = K_p$  and  $\mathbf{O}[p+1] = -k_{p+1} - 1$  such that we obtain

$$K_p = k_{p+1} = k_{p+1} + 1 - 1 = |\mathbf{O}[p+1]| - 1. \quad (6.23)$$

□

**Corollary 6.11.** *In the setting of Lemma 6.10 the number  $n_p$  of local trees of process  $p$  fulfills*

$$n_p = |\mathcal{O}[p+1]| - k_p = \begin{cases} |\mathcal{O}[p+1]| - \mathcal{O}[p] & \text{if } \mathcal{O}[p] \geq 0, \\ |\mathcal{O}[p+1]| - |\mathcal{O}[p]| + 1 & \text{else.} \end{cases} \quad (6.24)$$

*Proof.* This follows from the identity  $n_p = K_p - k_p + 1$ . □

Lemma 6.10 and Corollary 6.11 show that for valid partitions, the array  $\mathbf{0}$  carries the same information as the partition  $f$ .

### 6.2.3. Ghost trees

A valid partition provides information on the local trees of a process. These trees are all trees of which a forest has local elements. In many applications it is common to collect a layer of ghost (or halo) elements of the forest to support the exchange of data with neighboring processes. Since these ghost elements may be descendants of nonlocal trees, we store their trees as ghost trees. We want to confine this logic to the coarse mesh to be independent of a forest mesh, and thus we propose to store each nonlocal face-neighbor tree as a ghost tree. This means possibly storing more ghost trees than needed by a particular forest. However, this only affects the first and the last local tree of a process, which bounds the overhead. Since we restrict the neighbor information to face-neighbors, we also restrict ourselves to face-neighbor ghosts in this chapter. However, an extension to edge and vertex neighbor ghosts is planned for the future. These will prompt a somewhat more elaborate discussion, since an arbitrary number of trees can be neighbored across a vertex/edge. There exist known algorithms for quadrilaterals and hexahedra [70], which we believe can be modified to extend to simplices, prisms, and pyramids.

**Definition 6.12.** Let  $f$  be a valid partition of a coarse mesh. A **ghost tree** of a process  $p$  is any tree  $k$  such that

- $k \notin f(p)$ , and
- there exists a face-neighbor  $k'$  of  $k$  such that  $k' \in f(p)$ .

If a coarse mesh is partitioned according to  $f$ , then each process  $p$  will store its local trees and its ghost trees.

### 6.2.4. Computing the communication pattern

Suppose a coarse mesh is partitioned among the processes  $\{0, \dots, P-1\}$  according to a partition  $f$ . The input of the partition algorithm is this coarse mesh and a second partition  $f'$ , and the output is a coarse mesh that is partitioned according to the second partition.

Apart from these partitions being valid, no other restrictions are imposed on the partitions  $f$  and  $f'$ . Thus, we include the trivial case  $f = f'$  as well as extreme cases.

An example for such a case is an  $f$  that concentrates all trees on one process and an  $f'$  that assigns them to another (or distributes them evenly among all processes). For dynamic forest repartitioning it is not unusual for almost all of the elements to change their owner process [28, 100]. We expect similar behavior for the trees, especially when the number of trees is on the order of the number of processes or higher.

We suppose that in addition to its local trees and ghost trees, each process knows the complete partition tables  $f$  and  $f'$ , for example, in the form of offset arrays. The task is now for each process to identify the processes to which it needs to send local and ghost trees and then to execute the sending. A process also needs to identify the processes from which it receives local and ghost trees and to execute the receiving. We discuss here how each process can compute this information from the offset arrays without further communication.

It will become clear in section 6.2.5 that ghost trees need not yet be discussed at this point. Thus, it is sufficient to concentrate on the local trees for the time being.

### Ownership during partition

The fact that trees can be shared between multiple processes poses a challenge when repartitioning a coarse mesh. Suppose we have a process  $p$  and a tree  $k$  with  $k \in f'(p)$ , and  $k$  is a local tree for more than one process in the partition  $f$ . We do not want to send the tree multiple times, so how do we decide which process sends  $k$  to  $p$ ?

A simple solution would be that the process with the smallest index to which  $k$  is a local tree sends  $k$ . This process is unique and can be determined without communication. However, suppose that the two processes  $p$  and  $p-1$  share the tree  $k$  in the old partition, and  $p$  will also have this tree in the new partition. Then  $p-1$  would send the tree to  $p$  even though this message would not be needed.

We resolve this issue by only sending a local tree to a process  $p$  if this tree is not already local on  $p$ .

**Paradigm 6.13.** *When repartitioning with  $k \in f'(p)$ , the process that sends  $k$  to  $p$  is*

- $p$  if  $k$  already is a local tree of  $p$ , or else
- $q$  with  $q$  minimal such that  $k \in f(q)$ .

We acknowledge that sending from  $p$  to  $p$  in the first case is just a local data movement not involving communication.

**Definition 6.14.** When repartitioning, given a process  $p$  we define the sets  $S_p$  and  $R_p$  of processes to and from which  $p$  sends and receives local trees, respectively, and thus

$$S_p := \{ 0 \leq p' < P \mid p \text{ sends local trees to } p' \}, \quad (6.25a)$$

$$R_p := \{ 0 \leq p' < P \mid p \text{ receives local trees from } p' \}. \quad (6.25b)$$

Both sets may include the process  $p$  itself. Furthermore, we establish the notation for the smallest and largest ranks in these sets, understanding that they depend on  $p$ :

$$s_{\text{first}} := \min S_p, \quad s_{\text{last}} := \max S_p, \quad (6.26a)$$

$$r_{\text{first}} := \min R_p, \quad r_{\text{last}} := \max R_p. \quad (6.26b)$$

If  $S_p$  is empty, we set  $s_{\text{first}} = -1$  and  $s_{\text{last}} = -2$ , and likewise for  $R_p$ .  $S_p$  and  $R_p$  are uniquely determined by Paradigm 6.13.

### An example

We discuss a small example; see Figure 6.5. Here, we repartition a partitioned coarse mesh of five trees among three processes. The initial partition  $f$  is given by

$$0 = \{0, -2, 3, 5\} \quad (6.27)$$

and the new partition  $f'$  by

$$0' = \{0, -3, -4, 5\}. \quad (6.28)$$

Thus, initially tree 1 is shared by processes 0 and 1, while in the new partition, tree 2 is shared by processes 0 and 1, and tree 3 is shared by processes 2 and 3. We arrange the local trees that each process will send to every other process in a table, where the set in row  $i$ , column  $j$  is the set of local trees that process  $i$  sends to process  $j$ :

		0	1	2	
0		$\{0, 1\}$	$\emptyset$	$\emptyset$	(6.29)
1		$\{2\}$	$\{2\}$	$\emptyset$	
2		$\emptyset$	$\{3\}$	$\{3, 4\}$	

This leads to the following sets  $S_p$  and  $R_p$ :

$$S_0 = \{0\}, \quad R_0 = \{0, 1\}, \quad (6.30a)$$

$$S_1 = \{0, 1\}, \quad R_1 = \{1, 2\}, \quad (6.30b)$$

$$S_2 = \{1, 2\}, \quad R_2 = \{2\}. \quad (6.30c)$$

We see that process 1 keeps the tree 2 that is also needed by process 0. Thus, process 1 sends tree 2 to process 0. Process 0 also needs tree 1, which is local on process 1 in the old partition. But, since it is also local to process 0, process 1 does not send it.

### Determining $S_p$ and $R_p$

In this section we show that each process can compute the sets  $S_p$  and  $R_p$  from the offset array without further communication.




Figure 6.5.: A small example. The coarse mesh consists of five trees (numbered) and is partitioned among three processes (color coded). Left: the initial partition  $\mathcal{O}$ . Right: the new partition  $\mathcal{O}'$ . The colors of the trees encode the processes that have a tree as local tree. Process 0 is drawn in blue, process 1 in green, and process 2 in red. Initially, tree 1 is shared among processes 0 and 1, while in the new partition, tree 2 is shared among processes 0 and 1, and tree 3 is shared among processes 1 and 2. In (6.27)–(6.30) we list the sets  $\mathcal{O}$  and  $\mathcal{O}'$ , the trees that each process sends, and the sets  $S_p$  and  $R_p$ .

**Proposition 6.15.** *A process  $p$  can calculate the sets  $S_p$  and  $R_p$  without further communication from the offset arrays of the old and new partitions. Once the first and last elements of each set are known, process  $p$  can determine in constant time whether any given rank is in any of those sets.*

We split the proof into two parts. First, we discuss how a process can compute  $s_{\text{first}}$ ,  $s_{\text{last}}$ ,  $r_{\text{first}}$ , and  $r_{\text{last}}$ , and then show how it can decide for two processes  $\tilde{p}$  and  $q$  whether  $q \in S_{\tilde{p}}$ . In particular, we will apply this decision to each of the process numbers between the first and last elements of  $S_p$  and  $R_p$ .

We begin by determining  $s_{\text{first}}$  and  $s_{\text{last}}$  for  $S_p \neq \emptyset$ . For  $s_{\text{first}}$  we consider two cases. First, if the first local tree of  $p$  is not shared with a smaller rank, then  $s_{\text{first}}$  is the smallest process  $q$  that has this tree in the new partition and either is  $p$  itself or did not have it in the old one. We can find  $q$  with a binary search in the offset array.

Second, if the first tree of  $p$  is shared with a smaller rank, then  $p$  only sends it in the case that  $p$  keeps this tree in the new partition. Then  $s_{\text{first}} = p$ . Otherwise, we consider the second tree of  $p$  and proceed with a binary search as in the first case.

To compute  $s_{\text{last}}$ , we notice that among all ranks that have  $p$ 's old last tree in the new partition and did not already have it,  $s_{\text{last}}$  is the largest (except when  $p$  itself is this largest rank, in which case it certainly had the last tree). We can determine this rank with a binary search as well. If no such process exists, we proceed with the second-to-last tree of  $p$ , for which we know that such a process must exist.

**Remark 6.16.** The special case  $S_p = \emptyset$  occurs in the following situations:

1.  $p$  does not have any local trees.

2.  $p$  has one local tree that is shared with a smaller rank, and  $p$  does not have this tree as a local tree in the new partition.
3.  $p$  has two trees, the first of which case 2 holds for. The second (last) tree is shared with a set  $Q$  of bigger ranks, and there is no process  $q \notin Q$  that has this tree as a local tree in the new partition.

These conditions can be queried before computing  $s_{\text{first}}$  and  $s_{\text{last}}$ . To check condition 3, we need to perform one binary search, while we evaluate conditions 1 and 2 in constant time.

Similarly, to compute  $R_p$ , we first look at the smallest and largest elements of this set. These are the first and last processes from which  $p$  receives trees.  $r_{\text{first}}$  is the smallest rank that had  $p$ 's new first tree as a local tree in the old partition, or it is  $p$  itself if this tree was also a local tree on  $p$ . Also  $r_{\text{last}}$  is the smallest rank greater than or equal to  $r_{\text{first}}$  that had  $p$ 's new last local tree as a local tree in the old partition, or  $p$  itself. We can find both of these with a binary search in the offset array of the old partition.

**Remark 6.17.**  $R_p$  is empty if and only if  $p$  does not have any local trees in the new partition.

**Lemma 6.18.** *Given any two processes  $\tilde{p}$  and  $q$ , the process  $p$  can determine in constant time whether  $q \in S_{\tilde{p}}$ . Moreover,  $p$  can determine for a given tree  $k$  whether  $\tilde{p}$  sends  $k$  to  $q$ . In particular, this includes the cases  $\tilde{p} = p$  and  $q = p$ .*

*Proof.* Let  $\hat{k}_{\tilde{p}}$  be the first non-shared local tree of  $\tilde{p}$  in the old partition. If such a tree does not exist, then  $S_{\tilde{p}} = \emptyset$  or  $S_{\tilde{p}} = \{\tilde{p}\}$ . Let  $\hat{K}_{\tilde{p}}$  be the last local tree of  $\tilde{p}$  in the old partition if it is not the first local tree of  $q$  in the old partition, and let it be the second-to-last local tree otherwise. If such a second-to-last local tree does not exist, we conclude that  $\tilde{p}$  has only one tree in the old partition, and  $q$  also has this tree in the old partition. Thus  $q \notin S_{\tilde{p}}$ . Furthermore, let  $\hat{k}_q$  and  $\hat{K}_q$  be the first and last local trees of  $q$  in the new partition. We add 1 to  $\hat{k}_q$  if  $q$  sends its first local tree to itself, and this tree is also the new first local tree of  $q$ . We claim that  $q \in S_{\tilde{p}}$  if and only if all of the four inequalities

$$\hat{k}_{\tilde{p}} \leq \hat{K}_{\tilde{p}}, \quad \hat{k}_{\tilde{p}} \leq \hat{K}_q, \quad \hat{k}_q \leq \hat{K}_{\tilde{p}}, \quad \text{and} \quad \hat{k}_q \leq \hat{K}_q \quad (6.31)$$

hold. The only-if direction follows, since if  $\hat{k}_{\tilde{p}} > \hat{K}_{\tilde{p}}$ , then  $\tilde{p}$  does not have trees to send to  $q$ . If  $\hat{k}_{\tilde{p}} > \hat{K}_q$ , then the last new tree on  $q$  is smaller than the first old tree on  $\tilde{p}$ . If  $\hat{k}_q > \hat{K}_{\tilde{p}}$ , then the last tree that  $\tilde{p}$  could send is smaller than the first new local tree of  $p$ . Also if  $\hat{k}_q > \hat{K}_q$ , then  $q$  does not receive any trees from other processes. Thus,  $\tilde{p}$  cannot send trees to  $q$  if any of the four conditions is not fulfilled. The if-direction follows, since if all four conditions are fulfilled, there exists at least one tree  $k$  with

$$\hat{k}_{\tilde{p}} \leq k \leq \hat{K}_{\tilde{p}} \quad \text{and} \quad \hat{k}_q \leq k \leq \hat{K}_q. \quad (6.32)$$



Any tree with this property is sent from  $\tilde{p}$  to  $q$ . Process  $p$  can compute the four values  $\hat{k}_{\tilde{p}}$ ,  $\hat{K}_{\tilde{p}}$ ,  $\hat{k}_q$ , and  $\hat{K}_q$  from the partition offsets in constant time.  $\square$

**Remark 6.19.** Let  $p$  be a process that is not empty in the new partition. For symmetry reasons,  $R_p$  contains exactly those processes  $\tilde{p}$  with  $r_{\text{first}} \leq \tilde{p} \leq r_{\text{last}}$  and  $p \in S_{\tilde{p}}$ .

Thus, in order to compute  $S_p$ , we can compute  $s_{\text{first}}$  and  $s_{\text{last}}$  and then check for each rank  $q$  in between whether or not the conditions of Lemma 6.18 are fulfilled with  $\tilde{p} = p$ . For each process this check takes only constant run time.

Now, to compute  $R_p$  we can compute  $r_{\text{first}}$  and  $r_{\text{last}}$  and then check for each rank  $q$  in between whether or not  $p \in S_q$ .

These considerations complete the proof of Proposition 6.15.

### 6.2.5. Face information for ghost trees

We identify the following five different types of possible face connections in a coarse mesh:

1. Local tree to local tree.
2. Local tree to ghost tree.
3. Ghost tree to local tree.
4. Ghost tree to ghost tree.
5. Ghost tree to nonlocal and non-ghost tree.

There are several possible approaches to which of these face connections of a local coarse mesh we could actually store. As long as each face connection between any two neighbor trees is stored at least once globally, the information of the coarse mesh over all processes is complete, and a single process could reproduce all five types of face connection at any time, possibly using communication. Depending on which of these types we store, the pattern for sending and receiving ghost trees during repartitioning changes. Specifically, the tree that will become a ghost on the receiving process may be either a local tree or a ghost on the sending process.

When we use the maximum possible information of all five types of connections, we have the most data available and can minimize the communication required. In particular, from the nonlocal neighbors of a ghost and the partition table, a process can compute which other processes this ghost is also a ghost of and of which it is a local tree. With this information we can ensure that a ghost is sent only once and only from a process that also sends local trees to the receiving process.

The outline of the sending/receiving phase then looks like the following:

1. For each  $q \in S_p$ , send local trees that will be owned by  $q$  (following Paradigm 6.13).
2. Consider sending a neighbor of these trees to  $q$  if it will be a ghost on  $q$ . Send one of these neighbors if either  $p = q$  or both of the following conditions are fulfilled:

- $p$  is the smallest rank among those that consider sending this neighbor as a ghost, and
- $p \neq q$  and  $q$  does not consider sending this neighbor as a ghost to itself.

3. For each  $q \in R_p$ , receive the new local trees and ghosts from  $q$ .

In item 2 a process needs to know, given a ghost that is considered for sending to  $q$ , which other processes consider sending this ghost to  $q$ . This can be calculated without further communication from the face-neighbor information of the ghost. Since we know for each ghost the global index of each of its neighbors, we can check whether any of these neighbors is currently local on a different process  $\tilde{p}$  and will be sent to  $q$  by  $\tilde{p}$ . If so, we know that  $\tilde{p}$  considers sending this ghost to  $q$ .

Using this method, each local tree and ghost is sent only once to each receiver, and only those processes send ghosts that send local trees anyway, leading to minimal message numbers and message sizes. Storing less information would increase either the number of communicating processes or the amount of data that is communicated.


Supposing we did not store the face connection type 5, for ghost trees we would not have the information about to which nonlocal trees they connect. With this face information we could use a communication pattern such that each ghost is received only once by a process  $q$ , by sending the new ghost trees from a process that currently has it as a local tree (taking into account Paradigm 6.13). However, that process might not be an element of  $R_q$ , in which case additional processes would communicate.

If we stored only the local tree face information (types 1 and 2), then we would have minimal control over the ghost face connections. Nevertheless, we could define the partition algorithm by specifying that if a process  $p$  sends local trees to a process  $q$ , it will send all neighbors of these local trees as potential ghosts to  $q$ . The process  $q$  is then responsible for deleting those trees that it received more than once. With this method the number of communicating processes would be the same but the amount of data communicated would increase.

In Figure 6.6, we give an example comparing the three choices. To minimize the communication and overcome the need for postprocessing steps, we propose to store all five types of face connection.

### 6.3. Implementation

Let us begin by outlining the data structures for trees, ghosts, and the coarse mesh, and continue with a section on how to update the local tree and ghost indices. After this we present the partition algorithm to repartition a given coarse mesh according to a precalculated partition array. We emphasize that the coarse mesh data stores pure connectivity. In particular, it does *not* include the forest information, i.e., leaf elements and per-element payloads, which are managed by separate, existing algorithms.



	1, 2			1, 2, 3, 4			1, 2, 3, 4, 5		
p	0	1	2	0	1	2	0	1	2
0	0(1,2)	0(2)	—	0	0	(0)	0(1,2)	0	—
1	—	1(2)	2(0,1)	(1,2)	1(2)	2(1)	—	1(2)	2 (0,1)

Figure 6.6.: Repartitioning example of a coarse mesh showing the communication patterns controlled by the amount of face information available. Top: A coarse mesh of three trees is repartitioned. The numbers outside of the trees are their global indices. The numbers inside of each tree denote the processes that have this tree as a local tree. At first process 0 has tree 0, process 1 has trees 1 and 2, and process 2 has no local trees. After repartitioning process 0 has tree 0, process 1 has trees 0 and 1, and process 2 has tree 2 as local trees. Bottom: The table shows for each usage of face connection types which processes send which data. The row of process  $i$  shows in column  $j$  which local trees  $i$  sends to  $j$ , and—in parentheses—which ghosts it sends to  $j$ . Using face connection types 1–4 we use more communication partners (process 0 sends to process 2 and process 1 to process 0) than with all five types. Using types 1 and 2 only, duplicate data is sent (process 0 and process 1 both send the ghost tree 2 to process 1).

### 6.3.1. The coarse mesh data structure

Our data structure `cmesh` that describes a (partitioned) coarse mesh has the following entries:

- `0`: An array storing the current partition table; see Definition 6.9.
- `np`: The number of local trees on this process.
- `nghosts`: The number of ghost trees on this process.
- `trees`: A structure storing the local trees in order of their global indices.
- `ghosts`: A structure storing the ghost trees in no particular order.

We use 64-bit integers for global counts in `0` and use 32-bit signed integers for the local tree counts in `trees` and `ghosts`. This limits the number of trees per process to  $2^{31} - 1 \cong 2 \times 10^9$ . However, even with an overly optimistic memory usage of only 10 bytes per tree, storing that many trees would require about 18.6 GB of memory per process. Since on most distributed machines the memory per process is indeed much smaller, choosing 32-bit integers does not effectively limit the local number of trees. In presently unimaginable cases, we could still switch to 64-bit integers.

We call the index of a local tree inside the `trees` array the *local index* of this tree. Analogously, we call the index of a ghost in `ghosts` the *local index* of that ghost. On process  $p$ , we compute the global index  $k$  of a tree in `trees` from its local index  $\ell$  and compute the global index  $k_p$  of the first local tree and vice versa, since

$$k = k_p + \ell. \tag{6.33}$$

This allows us to address local trees with their local indices using 32-bit integers.

Each `tree` in the array `trees` stores the following data:

- `eclass`: The tree's shape as a small number (triangle, quadrilateral, etc.).
- `tree_to_tree`: An array storing the local tree and ghost neighbors along this tree's faces. See section 6.1.2 and the text below.
- `tree_to_face`: An array encoding for each face the face-neighbor's face number and the orientation of the face connection. See section 6.1.3.
- `tree_data`: A pointer to additional data that we store with the tree, for example, geometry information or boundary conditions defined by an application.

The  $i$ -th entry of `tree_to_tree` encodes the tree number of the face-neighbor at face  $i$  using an integer  $k$  with  $0 \leq k < n_p + n_{\text{ghosts}}$ . If  $k < n_p$ , the neighbor is the local tree with local index  $k$ . Otherwise, the neighbor is the ghost with local index  $k - n_p$ .

We do not allow a face to be connected to itself. Instead, we use such a connection in the face-neighbor array to indicate a domain boundary. However, a tree can be

connected to itself via two different faces. This allows for one-tree periodicity, as say in a 2D torus consisting of a single quadrilateral tree.

The `tree_data` field can contain arbitrary data. An application can use these, for example, to store higher-order geometry data per tree in order to account for curved boundaries of the coarse mesh. Refined elements in the forest can then be snapped to the curved boundaries by evaluating the `tree_data` field [27]. `tree_data` is partitioned to the processes together with the trees; thus possible duplicate copies of it can exist.

Each `ghost` in the array `ghosts` stores the following data:

- `Id`: The ghost's global tree index.
- `eclass`: The shape of the ghost tree.
- `tree_to_tree`: An array giving for each face the global number of its face-neighbor.
- `tree_to_face`: As above.

Since a ghost stores the global number of all its face-neighbor trees, we can locally compute all other processes that have this tree as a ghost by combining the information from `0` and `tree_to_tree`.

### 6.3.2. Updating local indices

After partitioning, the local indices of the trees and ghosts change. The new local indices of the local trees are determined by subtracting the global index of the first local tree from the global index of each local tree. The local indices of the ghosts are given by their positions in the data array.

Since the local indices change after repartitioning, we update the `tree_to_tree` entries of the local trees to store those new values. Because a neighbor of a tree can be either a local tree or a ghost on the previous owning process  $\tilde{p}$  and become either local or a ghost on the new owning process  $p$ , there are four cases that we shall consider.

We handle these four cases in two phases, the first phase being carried out on process  $\tilde{p}$  before the tree is sent to  $p$ . In this phase we change all neighbor entries of the trees that become local. The second phase executes on  $p$  after the tree has been received from  $\tilde{p}$ . At this point we change all neighbor entries belonging to trees that become ghosts.

In the first phase,  $\tilde{p}$  has information about the first local tree on  $\tilde{p}$  in the old partition, its global number being  $k_{\tilde{p}}$ . Via `0`' it also knows  $k_p^{\text{new}}$ , the global index of  $p$ 's first tree in the new partition. Given a local tree on  $\tilde{p}$  with local index  $\tilde{k}$  in the old partition, we compute its new local index  $k$  on  $p$  as

$$k = k_{\tilde{p}} + \tilde{k} - k_p^{\text{new}}, \quad (6.34)$$

which is its global index minus the global index of the new first local tree. Given a ghost  $g$  on  $\tilde{p}$  that will be a local tree on  $p$ , we compute its local tree number as

$$k = g.\text{Id} - k_p^{\text{new}}. \quad (6.35)$$

In the second phase,  $p$  has received all its new trees and ghosts and thus can give the new ghosts local indices to be stored in the `neighbors` fields of the trees. We do this by parsing its ghosts for each process  $\tilde{p} \in R_p$  (in ascending order) and incrementing a counter. For each ghost, we parse its neighbors for local trees, and for any of these we set the appropriate value in its `neighbors` field.

Note that these four cases apply in the special case  $\tilde{p} = p$  as well.

### 6.3.3. Partition\_cmesh: Algorithm 6.3.1

The input is a partitioned coarse mesh  $C$  and a new partition layout  $\mathcal{O}'$ , and the output is a new coarse mesh  $C'$  that carries the same information as  $C$  and is partitioned according to  $\mathcal{O}'$ .

This algorithm follows the method described in section 6.2.5 and is separated into two main phases, the `sending phase` and the `receiving phase`. In the former we iterate over each process  $q \in S_p$  and decide which local trees and ghosts we send to  $q$ . Before sending, we carry out phase one of the update of the local tree numbers. Subsequently, we receive all trees and ghosts from the processes in  $R_p$  and carry out phase two of the local index update.

In the `sending phase` we iterate over the trees that we send to  $q$ . For each of these trees we check for each neighbor (local tree and ghost) whether we send it to  $q$  as a ghost tree. This is the second item in the list of section 6.2.5. The function `Parse_neighbors` decides for a given local tree or ghost neighbor whether it is sent to  $q$  as a ghost.

## 6.4. Numerical results

The run time results that we present here have been obtained with version 0.2 of `t8code`<sup>1</sup> using the JUQUEEN supercomputer at Forschungszentrum Jülich, Germany. It is an IBM BlueGene/Q system with 28,672 nodes consisting of IBM PowerPC A2 processors at 1.6 GHZ with 16 GB RAM per node [73]. Each compute node has 16 cores and is capable of running up to 64 MPI processes using multithreading.

### 6.4.1. How to obtain example meshes

To measure the performance and memory consumption of the algorithms presented above, we would like to test the algorithms on coarse meshes that are too big to fit into the memory of a single process, which is 1 GB on JUQUEEN if we use 16

<sup>1</sup><https://github.com/cburstedde/t8code>

---

**Algorithm 6.3.1:** Partition\_cmesh(cmesh  $C$ , partition  $0'$ )

---

**Result:** A cmesh  $C'$  that consists of the same trees as  $C$  and is partitioned according to  $0'$ .

```
1  $p \leftarrow$  this process
2 From  $C.0$  and  $0'$  determine  $S_p$  and  $R_p$ . /* see section 6.2.4 */
   /* Sending phase */
3 for each  $q \in S_p$  do
4    $G \leftarrow \emptyset$  /* trees  $p$  sends as ghosts to  $q$  */
5    $s \leftarrow$  first local tree to send to  $q$ .
6    $e \leftarrow$  last local tree to send to  $q$ .
7    $T \leftarrow \{C.trees[s], \dots, C.trees[e]\}$  /* local trees  $p$  sends to  $q$  */
8   for  $k \in T$  do
9     Parse_neighbors( $C, k, q, G, 0', s, e$ )
10    update_tree_ids_phase1( $T \cup G$ ) /* see equations (6.34) and (6.35) */
11    Send  $T \cup G$  to process  $q$ 
   /* Receiving phase */
12 for each  $q \in R_p$  do
13   Receive  $T[q] \cup G[q]$  from process  $q$ 
14  $C'.trees \leftarrow \bigcup_{R_p} T[q]$  /* new array of local trees */
15  $C'.ghosts \leftarrow \bigcup_{R_p} G[q]$  /* new array of ghost trees */
16 update_tree_ids_phase2( $C'.ghosts$ )
17  $C'.0 \leftarrow 0'$ 
18 return  $C'$ 
```

---

```
   /* decide which neighbors of  $k$  to send as a ghost to  $q$  */
1 Function Parse_neighbors(cmesh  $C$ , tree  $k$ , process  $q$ , ghosts  $G$ , partition  $0'$ ,
   tree_indices  $s, e$ )
2 for  $u \in k.tree\_to\_tree \setminus \{s, \dots, e\}$  do
3   if  $0 \leq u < n_p$  and Send_ghost( $C, ghost(u), q, 0'$ ) then
4     if  $u + k_p \notin f'(q)$  then
5        $G \leftarrow G \cup \{ghost(u)\}$  /* local tree  $u$  becomes ghost of  $q$  */
6     else /*  $n_p \leq u$  */
7        $g \leftarrow C.ghosts[u - n_p]$ 
8       if  $g.Id \notin f(q)$  and Send_ghost( $C, g, q, 0'$ ) then
9          $G \leftarrow G \cup \{g\}$  /*  $g$  is a ghost of  $q$  */
```

---

```
   /* Subroutine to decide whether to send a ghost or not */
1 Function Send_ghost(cmesh  $C$ , ghost  $g$ , process  $q$ , partition  $0'$ )
2  $S \leftarrow \emptyset$ 
3 for  $u \in g.tree\_to\_tree$  do
4   for  $q'$  with  $u$  is a local tree of  $q'$  do
5     if  $q'$  sends  $u$  to  $q$  then /* See Lemma 6.18 */
6        $S \leftarrow S \cup \{q'\}$ 
7 if  $q \notin S$  and  $p = \min S$  then
8   return true /*  $p$  is the smallest rank sending trees to  $q$  */
9 else
10  return false
```

---

MPI ranks per node. We consider the following three approaches to construct such meshes:

1. Use an external parallel mesh generator.
2. Use a serial mesh generator on a large-memory machine, transfer the coarse mesh to the parallel machine’s file system, and read it using (parallel) file I/O.
3. Create a large coarse mesh by forming the disjoint union of smaller coarse meshes over the individual processes.

Due to a lack of availability of parallel open source mesh generators, we restrict ourselves to the second and third methods. These have the advantage of being started with initial coarse meshes that fit into a single process’s memory such that we can work with serial mesh generating software. In particular, we use `gmsh`, `TetGen`, and `Triangle` [57, 127, 128].

The third method is especially well suited for weak scaling studies, since the individual small coarse meshes can be created programmatically and communication-free on each process. They may be of the same size or different sizes among the processes.

We discuss two examples below. In the first we examine purely the coarse mesh partitioning without regard for a forest and its elements (using hexahedral meshes), and in the second we drive the coarse mesh partitioning by a dynamically changing forest of elements (using tetrahedral meshes). The latter example produces shared trees and thus fully executes the algorithmic ideas put forward above.

#### 6.4.2. Disjoint bricks

In our first example we conduct both strong and weak scaling studies of coarse mesh partitioning and test the maximal number of (hexahedral) trees that we can support before running out of memory. For our weak scaling results, we keep the same per-process number of trees while increasing the total number of processes, which we achieve by constructing an  $n_x \times n_y \times n_z$  brick of trees on each process using three constant parameters  $n_x$ ,  $n_y$ , and  $n_z$ . We repartition this coarse mesh once, by the rule that each rank  $p$  sends 43% of its local trees to the rank  $p + 1$  (except the biggest rank  $P - 1$ , which keeps all its local trees). We choose this odd percentage to create nontrivial boundaries between the regions of trees to keep and trees to send. See Figure 6.7 for a depiction of the partitioned coarse mesh on six processes. The local bricks are created locally as `p4est` connectivities with `p4est_connectivity_new_brick` and are then reinterpreted in parallel as a distributed coarse mesh data structure.

We perform strong and weak scaling studies on up to 917,504 MPI ranks and display our results in Figures 6.8 and 6.9 and Table 6.2. We show the results of one study with 16 MPI ranks per compute node, thus 1 GB available memory per process, and one with 32 MPI ranks per compute node, leaving half of the memory per process. In both cases we measure run times for different mesh sizes per process.






Figure 6.7.: The structure of the coarse mesh that we use to measure the maximum possible mesh sizes and scalability for an example with six processes. Before partitioning, the coarse mesh local to each process is created as one  $n_x \times n_y \times n_z$  block of hexahedral trees. We repartition the mesh such that each process sends 43% of its local trees to the next process. The picture shows the resulting partitioned coarse mesh with parameters  $n_x = 10$ ,  $n_y = 18$ , and  $n_z = 8$  and color coded MPI rank.

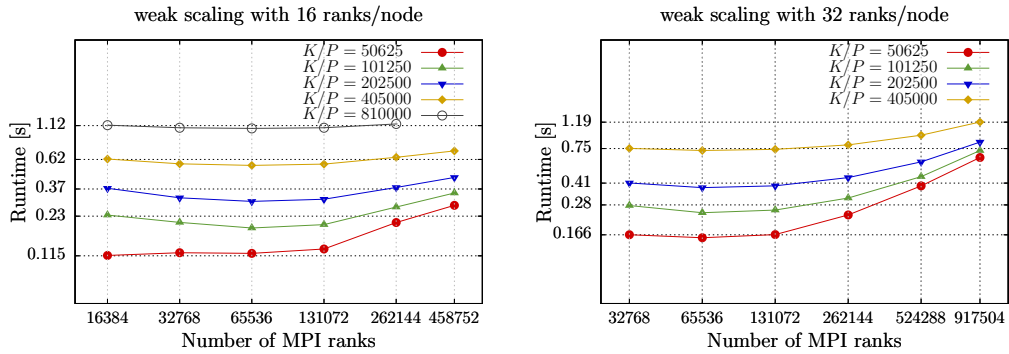


Figure 6.8.: Weak scaling of `Partition_cmesh` with disjoint bricks. Left: 16 ranks per node. Right: 32 ranks per node. We show the run times for the baseline on the y-axis and provide graphs for different ratios between total coarse cells  $K$  and MPI processes  $P$ . On the left-hand side the time for the largest 458,752 process run is 0.72 seconds; on the right-hand side the time for the largest 917,504 process run is 1.19 seconds. We obtain efficiencies of  $0.62/0.72 = 86\%$  and  $0.75/1.19 = 63\%$  compared to the baselines of 16,384/32,768 MPI ranks, respectively (yellow lines). The time for the 262,144 process run with 810e3 trees per process (black line) increases from 1.12 to 1.15 seconds, which translates into a weak scaling efficiency of 97.4%. (See online version for color.)

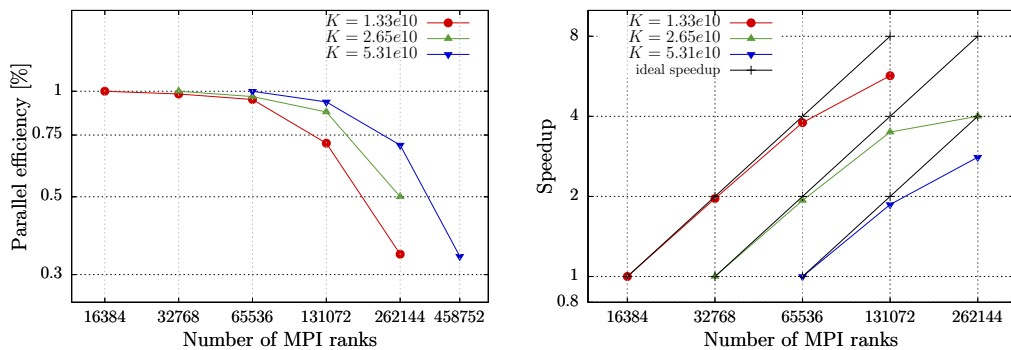


Figure 6.9.: Strong scaling of `Partition_cmesh` for the disjoint bricks example on JUQUEEN with 16 ranks per compute node, for three runs with 13.3e9, 26.5e9, and 53.1e9 trees. We show the parallel efficiency on the left and the speedup on the right. The absolute run times for 262,144 processes are 0.21, 0.27, and 0.38 seconds.

Run time tests for `Partition_cmesh`

131,072 MPI ranks (16 ranks per node)					
Mesh size	Per rank	Trees (ghosts) sent		Time [s]	Factor
6.635e9	50,625	21,767	(3,414)	0.13	–
13.27e9	101,250	43,536	(5,504)	0.20	1.53
26.54e9	202,500	87,074	(6,607)	0.31	1.56
53.08e9	405,000	174,149	(11,381)	0.57	1.85
106.2e9	810,000	348,297	(22,335)	1.08	1.89
917,504 MPI ranks (32 ranks per node)					
Mesh size	Per rank	Trees (ghosts) sent		Time [s]	Factor
46.45e9	50,625	21,768	(3,413)	0.64	–
92.90e9	101,250	43,537	(5,504)	0.72	1.13
185.8e9	202,500	87,075	(6,607)	0.84	1.12
371.6e9	405,000	174,150	(11,383)	1.19	1.42

Table 6.2.: The run times of `Partition_cmesh` for 131,072 processes with 16 processes per node (top) and 917,504 processes with 32 processes per node (bottom). The largest coarse mesh that we created during the tests has 371e9 trees. In the middle column we list the average number of trees (ghosts) that each process sends to another process. The last column is the quotient of the current run time divided by the previous run time. Since we double the mesh size in each step, we expect an increase in run time of a factor of 2, which hints at parallel overhead becoming negligible in the limit of many trees per process.

#MPI ranks	#trees	Run time [s]
1,024	4,096	0.00136
1,024	8,192	0.00149
1,024	16,384	0.00142
64	105	0.00122
32	105	0.00789
64	3,200	0.000293
64	19,200	0.000865

Table 6.3.: Run times for `Partition_cmesh` for relatively small coarse meshes. The bottom two rows of the table was not computed on JUQUEEN but on a local institute cluster of 78 nodes with 8 Intel Xeon CPU E5-2650 v2 @ 2.60GHz each.

We observe that even for the biggest meshes of 405e3 and 810e3 trees per process the absolute run times of partition are below 1.2 seconds. Furthermore, we measure a weak scaling efficiency of 97.4% for the 810e3 mesh on 262,144 processes and 86.2% for the 405e3 mesh on 458,752 processes. The biggest mesh that we created is partitioned between 917,504 processes and uses 405e3 trees per process for a total of over 371e9 trees.

We notice a drop off in the scaling behavior when the number of trees per process is about 100e3 and smaller. At this stage the time for local computation is small relative to the time for communication. Since in many cases the number of trees per process will likely be even smaller, we add Table 6.3. It documents running `Partition_cmesh` with small coarse meshes, using a number of trees on the order of the number of processes. These tests show that for such small meshes the run times are on the order of milliseconds. Hence, for small meshes there is no disadvantage in using a partitioned coarse mesh over a replicated one (i.e., each process holding a full copy).

### 6.4.3. An example with a forest

In this example we partition a tetrahedral coarse mesh according to a parallel forest of fine elements. While we pushed the maximum number of trees in the previous example, we now consider mesh sizes that occur in more common usage scenarios.

When simulating shock waves or two-phase flows, there is often an interface along which a finer mesh resolution is desired in order to minimize computational errors. Motivated by this example, we create the forest mesh as an initial uniform refinement of the coarse mesh with a specified level  $\ell$  and refine it in a band along an interface defined by a plane in  $\mathbb{R}^3$  up to a maximum refinement level  $\ell + k$ . As the refinement rule we use 1:8 red refinement [16] together with the tetrahedral Morton SFC [29]. We move the interface through the domain with a constant velocity. Thus, in each time step the mesh is refined and coarsened, and therefore we repartition it to maintain an optimal load balance. We measure run times for both coarse mesh and forest mesh partitioning for three time steps.

Our coarse mesh consists of tetrahedral trees modeling a brick with spherical holes in it. To be more precise, the brick is built out of  $n_x \times n_y \times n_z$  tetrahedralized unit cubes, and each of those has one spherical hole in it; see Figures 6.10 and 6.11 for a small example mesh.

We create the mesh in serial using the generator `gmsh` [57]. We read the whole file on a single process, and thus use a local machine with 1 terabyte memory for preprocessing. On this machine we partition the coarse mesh to several hundred processes and write one file for each partition. This data is then transferred to the supercomputer. The actual computation consists of reading the coarse mesh from files, creating the forest on it, and partitioning the forest and the coarse mesh simultaneously. To optimize memory while loading the coarse mesh, we open at most one partition file per compute node.

The coarse mesh that we use in these tests has parameters  $n_x = 26, n_y = 24, n_z = 18$  and thus 11,232 unit cubes. Each cube is tetrahedralized with about 34,150 tetrahedra, and the whole mesh consists of 383,559,464 trees. In the first test, we create a forest of uniform level 1 and maximal refinement level 2, and in the second, we create a forest of uniform level 2 and maximal refinement level 3. The forest mesh in the first test consists of approximately 2.6e9 elements. In the second test, we also use a broader band and obtain a forest mesh of 25e9 tetrahedra.

In Table 6.4 we show the run time results and further statistics for coarse mesh partitioning and in Table 6.5 show results for forest partitioning. We observe that the run time for `Partition_cmesh` is between 0.10 and 0.13 seconds, and that about 88% or 98%, respectively, of all processes share local trees with other processes. The run times for forest partition are below 0.22 seconds for the first example and below 0.65 seconds for the second example.

We also run a third test on 458,752 MPI ranks and refine the forest to a maximum level of four. Here, the forest mesh has 167e9 tetrahedra, which we partition in under 0.6 seconds. The coarse mesh partition routine runs in about 0.2 seconds. Approximately 60% of the processes have a shared tree in the coarse mesh. In Table 6.6 we show the results from this test.




Figure 6.10.: The coarse mesh connectivity that we use for the partition tests motivated by an adapted forest. It consists of  $n_x \times n_y \times n_z$  cubes, with each cube having one spherical hole. For this picture we use  $n_x = 4$ ,  $n_y = 3$ ,  $n_z = 2$ , and each cube is triangulated with approximately 7,575 tetrahedra. For illustration purposes we show some parts of the mesh as opaque and other parts as invisible.




Figure 6.11.: An illustration of the band of finer forest mesh elements in the example. The region of finer mesh elements moves through the mesh in each time step. From left to right we see  $t = 1$ ,  $t = 2$ , and  $t = 3$ . In this illustration, elements of refinement level 1 are blue and elements of refinement level 2 are red.

t	Trees (ghosts) sent	Data sent [MiB]	$ S_p $	Shared trees	Run time [s]
1	25,117 (11,948)	4.95	2.27	7,178	0.103
2	34,860 (16,854)	6.88	2.75	7,176	0.110
3	36,386 (17,568)	7.18	2.82	7,182	0.112
1	39,026 (18,334)	7.67	2.97	8,096	0.128
2	38,990 (18,268)	7.66	2.95	8,085	0.129
3	38,942 (18,074)	7.64	2.93	8,085	0.128

Table 6.4.: `Partition_cmesh` with a coarse mesh of 324,766,336 tetrahedral trees on 8,192 MPI ranks. We measure the duration of mesh repartitioning for three time steps. For each one, we show process average values of the number of trees (ghosts) and the total number of bytes that each process sends to other processes. The average number of other processes to which a process sends ( $|S_p|$ ) is below three in each test. We also provide the total number of shared trees in the mesh, where 8,191 is the maximum possible value.

t	Mesh size	Elements sent	Data sent [MiB]	Run time [s]
1	2,622,283,453	203,858	3.49	0.215
2	2,623,842,241	281,254	4.82	0.215
3	2,626,216,984	293,387	5.03	0.214
1	25,155,319,545	3,013,230	46.6	0.642
2	25,285,522,233	3,008,800	46.5	0.640
3	25,426,331,342	2,991,990	46.2	0.645

Table 6.5.: Forest mesh partition on 8,192 MPI ranks. For the same example as in Table 6.4 we display statistics and run times for the forest mesh partition. We show the total number of tetrahedral elements and the average count of elements and bytes that each process sends to other processes (their count is the same as in Table 6.4).

t	Trees (ghosts) sent	Data sent [MiB]	$ S_p $	Shared trees	Run time [s]
1	704 (2,444)	0.267	2.99	280,339	0.207
2	707 (2,456)	0.269	3.00	281,694	0.204
3	708 (2,458)	0.269	3.00	281,900	0.204

t	Mesh size	Elements sent	Data sent [MiB]	Run time [s]
1	167,625,595,829	362,863	5.55	0.522
2	167,709,936,554	364,778	5.58	0.578
3	167,841,392,949	365,322	5.59	0.567

Table 6.6.: Coarse and forest mesh partitions on 458,752 MPI ranks. Run times for coarse mesh (top) and forest mesh (bottom) partition for the brick with holes on 458,752 MPI ranks. The setting and the coarse mesh are the same as in Table 6.4 except that for the forest we use an initial uniform level three refinement with a maximum level of four.

## 6.5. Conclusion

In this chapter we propose an algorithm that executes dynamic and in-core coarse mesh partitioning. In the context of tree-based adaptive mesh refinement (AMR), the coarse mesh defines the connectivity of tree roots, which is used in all neighbor query operations between elements. This development is motivated by simulation problems on complex domains that require large input meshes. Without partitioning of the tree meta data, we will run out of memory around one million trees, and with static or out-of-core partitioning, we might not have the flexibility to transfer the tree meta data as required by the change in process ownership of the trees' elements, which occurs in every AMR cycle. With the approach presented here, this can be performed with run times that are significantly smaller than those for partitioning the elements, even considering that SFC methods for the latter are exceptionally fast in absolute terms. Thus, we add little to the run time of all AMR operations combined.

Our algorithm guarantees that each process can provide the tree meta data for each of its fine mesh elements that are themselves distributed using a SFC. We handle the communication without handshaking and develop a communication pattern that minimizes data movement. This pattern is calculated by each process individually, reusing information that is already present.

Our implementation scales up to  $917e3$  MPI processes and up to  $810e3$  trees per process, where the largest test case consists of  $371e9$  trees. What remains to be done is extending the partitioning of ghost trees to edge and corner neighbors, since only face-neighbor ghost trees are presently handled. It appears that the structure of the algorithm will allow this with little modification.



## 7. Ghost

For many algorithms we need to know information about all neighbors of a forest mesh element. A typical example is if an application needs to calculate the fluxes in a finite volume solver or to compute integrals in a finite element setting; see for example [20, 119, 149] and our own discussion in Chapter 9. Other examples include a refinement criterion that takes the size of neighboring elements into account or the gradient of an approximated function.

If the forest mesh is partitioned among multiple processes, then a neighbor element of a leaf element owned by process  $p$  may be owned by a different process  $q \neq p$ . We call this neighbor a **ghost** element of  $p$ .

In this section we describe how we create a layer of ghost elements for a partitioned forest. Thus, each process obtains information about all of its ghost elements.

As before, we restrict the algorithm to face-neighbors.

**Definition 7.1.** A **ghost element** (or **ghost** for short) of a process  $p$  in a forest  $\mathcal{F}$  is a leaf element  $G$  of a process  $q \neq p$ , such that there exists a face-neighbor  $E$  of  $G$  that is a local leaf element of  $p$ .

**Definition 7.2.** We call a local element  $E$  **boundary element** if it has at least one face-neighbor that is a ghost element. The **remote processes** of  $E$  are all processes  $q \neq p$  that own ghost elements of  $E$ . The union of all remote processes of all local elements of  $p$  are the remote processes of  $p$ .

**Definition 7.3.** By  $R_p^q$  we denote the set of boundary elements of process  $p$  that have process  $q$  as a remote process.

Throughout this section we assume that we can access the coarse mesh information of each neighbor tree of a local tree. We can do this, since the coarse mesh is either replicated or stores a layer of ghost trees; see Section 6.2.3.

In the following, we need the definition of a  $(2 : 1)$ -balanced forest.

**Definition 7.4.** We call a forest  $\mathcal{F}$  **balanced** if each pair  $E, E'$  of face-neighboring leaf elements of  $\mathcal{F}$  satisfies

$$\ell(E) - 1 \leq \ell(E') \leq \ell(E) + 1. \quad (7.1)$$

Here,  $E$  and  $E'$  may belong to different processes. Thus, any two face-neighbors differ by at most one in their refinement levels. If the condition is not fulfilled, we say that  $\mathcal{F}$  is unbalanced.

**Remark 7.5.** In some publications *graded* is used instead of *balanced* [41, 104].

We describe a basic version of a **Balance** algorithm that transforms an unbalanced forest into a balanced one by successively refining elements in Chapter 8.

In this chapter we discuss three steps of implementing **Ghost**, which we denote by **Ghost\_v1**, **Ghost\_v2**, and **Ghost\_v3**. **Ghost\_v1** is a relatively straight-forward version that only works with balanced forests. **Ghost\_v2** is a more sophisticated version that works on arbitrary forests, and we optimize its runtime in **Ghost\_v3**. For the first two versions we orient ourselves to the **p4est** implementations for quadrilateral/hexahedral meshes [34, 70]. However, we discuss them in our new, element-type independent framework, which not only extends to triangular/tetrahedral meshes, but also to hybrid meshes consisting of different element types. To this end, we out-source all operations that require specific knowledge of the element type as low-level functions (c.f. Section 2.5.3). Changing the element type is equivalent to changing the low-level implementation. In this chapter, we discuss implementations for line, quadrilateral, and hexahedral elements with the Morton index, as well as for triangular and tetrahedral elements with the TM-index. This implicitly gives us an implementation of prism elements, since we can model these as the cross product of a line and a triangle; see [81].

The basic idea of **Ghost\_v1** and **Ghost\_v2** is to first identify all boundary elements and their remote processes, thus building the sets  $R_p^q$  and identifying the non-empty ones. In a second step, each process  $p$  sends all elements in  $R_p^q$  to  $q$ .

In the first step we iterate over all local leaves and for each over all of its faces. We then have to decide for each face  $F$  of a leaf  $E$  which processes own leaves that touch this face. The difference between **Ghost\_v1** and **Ghost\_v2** lies in this decision process.

A new improvement, **Ghost\_v3**, replaces the iteration over all leaves with a top-down search. With this approach, we exclude portions of the mesh from the iteration, if they lie entirely within a process's domain, which improves the overall runtime by at least one order of magnitude. In **p4est** the runtime is optimized by performing a so called  $(3 \times 3)$ -neighborhood check of an element [70]. For a local hexahedral (or quadrilateral, in 2D) element, we check whether all same-level face- (or edge-/vertex-) neighbors are also process local and if so, the element is excluded from the iteration. Since this check makes explicit use of the Morton code and its properties, it is difficult to generalize for general element types. Therefore, we choose a different ansatz by exploiting the top-down search, which directly leads to an element-type independent algorithm.

Note that for unbalanced forests the number of neighbors of an element  $E$  that are ghosts can be arbitrarily large and is only bounded by the number of elements at maximum refinement level that can touch the faces of  $E$ . Therefore, also the number of remote processes is unbounded as well.

## 7.1. Element face-neighbors

An important part of `Ghost` is to construct the same-level face-neighbor of a given element  $E$  across a face  $f$ . As long as such a face-neighbor is inside the same tree as  $E$ , this problem is solved by the corresponding low-level function `t8_element_face_neighbor_inside`. We describe its version for the TM-index in Algorithm 4.3.6; see [34] for an implementation for the hexahedral Morton index.

The challenging part is to find element face-neighbors across tree boundaries. This is particularly demanding for hybrid meshes since multiple types of trees exist in the same forest; this is for example the case if a hexahedron tree is neighbor of a prism tree. We thus aim for a general type-independent algorithm to compute face-neighbors across tree boundaries.

To detect whether an element's face  $f$  is also a tree boundary, we assume the existence of a low-level algorithm `element_neighbor_inside_root`<sup>1</sup> that returns `true` if and only if the face  $f$  is an inner face and thus not on the tree boundary. For the TM-index we can derive such an algorithm from the methods that we present in Section 4.3.4.

**Remark 7.6.** In the following we will denote all functions that are part of the `t8code` low-level API with a `t8_element` prefix. See Appendix A for a complete list.

The core idea for the face-neighbor algorithm across tree boundaries is to explicitly build the face as a lower dimensional element.

Our notation convention is to use capital letters  $(T, E, F, G)$  for entities, such as trees, elements and faces, and to use lower case for indices  $(t, e, f, g)$  of those entities.

We are concerned with the following issue:

*Given a  $d$ -dimensional element  $E$  in a tree  $K$  and a face  $F$  of  $E$  of which we know that it is a subface of a face  $G$  of  $K$ , construct the same-level face-neighbor element  $E'$  of  $E$  across  $F$ .*

We display this situation for a quadrilateral-triangle tree connection in Figure 7.1.

**Remark 7.7.** We emphasize that neither the original element  $E$  nor its face-neighbor  $E'$  need to be leaves in the forest. Even if  $E$  is a leaf in the forest,  $E'$  does not necessarily have to be a leaf, since  $E'$  has the same refinement level as  $E$ .

Since the trees can be rotated against each other, their coordinate systems may not be aligned. We must properly transform the  $(d - 1)$ -dimensional coordinates of the faces between the two systems.

In order to compute the face-neighbor  $E'$  we consider the corresponding face  $G$  of the tree  $K$  as a  $(d - 1)$ -dimensional root tree. We then explicitly construct the

---

<sup>1</sup>In `t8code` this function is part of `t8_element_face_neighbor_inside`; see Appendix A




Figure 7.1.: We show a tree  $K$ , an element  $E$ , and a face  $F$  of  $E$  that is a subface of a tree face. The task is to find the face-neighbor element  $E'$ . A subtask is to identify the tree faces  $G$  and  $G'$ , and the face  $F'$ .

face  $F$  of  $E$  as a  $(d - 1)$ -dimensional element descendant of  $G$ . The next step is to transform the coordinates of  $F$  accordingly to build the  $(d - 1)$ -dimensional element  $F'$  as a descendant of  $G'$ , the face's root tree. In a final step, we construct the element  $E'$  from the given face element  $F'$ .

We thus identify four major substeps in the computation of face-neighbors across tree boundaries:

- (i) Identify the face number  $g$  of the tree face.
- (ii) Construct the  $(d - 1)$ -dimensional face element  $F$ .
- (iii) Transform the coordinate system of  $F$  to obtain the neighbor face element  $F'$ .
- (iv) Extrude  $F'$  to the  $d$ -dimensional neighbor element  $E'$ .

We show these steps in Algorithm 7.1.1 and describe their details in the following sections. See Figure 7.2 for an illustration of steps (ii), (iii), and (iv).

**Remark 7.8.** We deliberately choose this method of using lower dimensional entities over directly transforming the tree coordinates from one tree to the other—as it is done for example in [34]—since our approach allows for maximum flexibility of the implementations of the different element and SFC types. This holds since all intermediate operations are either local to one element type or only change the dimension (i.e. hexahedra to quadrilaterals, tetrahedra to triangles, and back). Therefore, even if, for example, a hexahedron tree is neighbor to a prism tree, no function in the implementation of the hexahedral elements relies on knowledge about the implementation of the prism elements. Hence, it is possible to exchange different implementations of SFC of one element type without changing the others.




Figure 7.2.: Two hexahedral elements that are face-neighbors across tree boundaries (exploded view). Constructing the face-neighbor  $E'$  of  $E$  across the face amounts to computing its anchor node (black) from the anchor node of  $E$  and the coarse mesh connectivity information about the tree connection. Here, the coordinate systems of the two trees are rotated against each other. Top right: In step (ii) we construct the face element  $F$  from the element  $E$ . The coordinate system of the face root is inferred from that of the left tree. Bottom left: In step (iii) we transform the face element  $F$  to the neighbor face element  $F'$ . Bottom right: In the last step (iv) we extrude the face-neighbor  $E'$  from the face element  $F'$ .

---

**Algorithm 7.1.1:** `t8_forest_face_neighbor` (Forest  $\mathcal{F}$ , tree  $T$ , element  $E$ , face number  $f$ )

---

```

Result: The same-level face-neighbor  $E'$  of  $E$  across  $f$ .
1 if element_neighbor_inside_root ( $E, f$ ) then
2    $E' \leftarrow$  t8_element_face_neighbor_inside ( $E, f$ )
3   return  $E'$ 
4  $g \leftarrow$  t8_element_tree_face ( $E, f$ )           /* (i) The face number of  $G$  */
5  $F \leftarrow$  t8_element_boundary_face ( $E, f$ )     /* (ii) Construct the face element */
6  $o \leftarrow$  face_orientation ( $\mathcal{F}, T, g$ )       /* The orientation of the tree face connection */
7  $F' \leftarrow$  t8_element_transform_face ( $F, o$ )   /* (iii) Obtain the neighbor face element */
8  $g' \leftarrow$  tree_neighbor_face ( $\mathcal{F}, T, g$ )     /* The face number of  $G'$  */
9  $E' \leftarrow$  t8_element_extrude_face ( $F', g'$ )  /* (iv) Build the neighbor from  $F'$  */
10 return  $E'$ 

```

---

### 7.1.1. (i) Identifying the tree face

The first subproblem is to identify the tree face  $G$ , respectively its face index  $g$ , from  $E$ ,  $f$ , and  $K$ . For this task we define a new low-level function:

- `t8_element_tree_face` (element  $E$ , face\_index  $f$ )

If  $f$  is a subface of a tree face, return the face index  $g$  of this root tree face.

The function `t8_element_tree_face` returns the root tree face index of the root face of an element's face  $f$ , provided this face is a boundary face. Its return value thus depends on the enumeration of the faces of an element in relation to the faces of its root tree.

For lines, quadrilaterals and hexahedra with the Morton index, the root tree face indices are the same as the element's face indices [34] and thus `t8_element_tree_face` always returns  $f$ .

For simplices with the TM index, the enumeration of their faces depends on their simplex type. Face number  $i$  refers to the unique face that does not contain the vertex  $\vec{x}_i$ . We show the vertices of the different types in Figure 4.2.

We observe that for triangles of type 0, the face number is the same as the face number of the root tree (since triangles of type 0 are scaled copies of the root tree). Triangles of type 1 cannot lie on the boundary of the root tree and thus we never call `t8_element_tree_face` with a type 1 triangle.

For tetrahedra of type 0 the same reasoning holds as for type 0 triangles, `t8_element_tree_face` returns  $f$ . Tetrahedra of type 3 cannot lie on the boundary of the root tree. For each of the remaining four types there is exactly one face that can lie on the root tree boundary. Face 0 of type 1 tetrahedra is a descendant of the root face 0; face 2 of type 2 tetrahedra is a descendant of the root face 1; face 1 of type 4 tetrahedra is a descendant of the root face 2. Finally, face 3 of type 5 tetrahedra is a descendant of the root face 3. We list these indices in Table 7.1.

Note that for face indices  $f$  of faces that cannot lie on the root boundary, the return value of `t8_element_tree_face` is undefined. This behavior is legal, since we ensure in Algorithm 7.1.1 that the function is only called if the face  $f$  does lie on the root boundary. We do so by calling `element_neighbor_inside_root` beforehand which queries whether the face-neighbor across  $f$  is in the root tree or not.

### 7.1.2. (ii) Constructing the face element

As a next step, we build the face  $F$  as a  $(d - 1)$ -dimensional element. We do this via the low-level function

- `t8_element_boundary_face` (element  $E$ , face\_index  $f$ )

Return the  $(d - 1)$ -dimensional face element  $F$  corresponding to the face index  $f$ .

Thus, the lower dimensional face element  $F$  has to be created from  $E$ . For the Morton index this is equivalent to computing the coordinates of its anchor node

Tetrahedron					
type( $T$ )	$f$	$g$	type( $T$ )	$f$	$g$
0	$i$	$i$	3	-	-
1	0	0	4	1	2
2	2	1	5	3	3

Table 7.1.:  $g = \text{t8\_element\_tree\_face}(T, f)$  for a tetrahedron  $T$  and a face  $f$  of  $T$  that lies on a tree face. Depending on  $T$ 's type, all, exactly one, or none of its faces can be a subface of a face of the root tetrahedron tree. We show the tetrahedron's face number  $f$  and the corresponding face number  $g$  in the root tetrahedron. Type 3 tetrahedra can never have a subface of the root tetrahedron as face. For type 0 tetrahedra, each of their faces can be a subface of the root tetrahedron's face with the same index.




Figure 7.3.: Constructing the face element  $F$  to an element  $E$  at a tree face  $G$ . We can interpret the face of the 3D tree  $K$  as a 2D tree  $G$ . The face  $F$  of  $E$  is an element in this tree.

and additionally its type for simplices. Hereby we interpret the tree face  $G$  as a  $(d-1)$ -dimensional root tree of which  $F$  is a descendant element. See also Figure 7.3.

**Remark 7.9.** Since we construct a lower-dimensional element as the face of a higher-dimensional one, there are two conditions that need to be satisfied for the implementations of the two element types involved.

1. The refinement pattern of a face of the higher dimensional elements must match the lower dimensional refinement pattern.
2. The maximum possible refinement level of higher dimensional elements must not exceed the one of the lower dimensional elements.

If one, or both, of these conditions are not fulfilled, then there exist faces of the higher dimensional elements for which an interpretation as a lower dimensional element is not possible. For Morton-type SFCs these conditions are naturally fulfilled.

Quad		Hexahedron			
$f$	$F.x$	$f$	$(F.x, F.y)$	$f$	$(F.x, F.y)$
0	$Q.y$	0	$(Q.y, Q.z)$	3	$(Q.x, Q.z)$
1	$Q.y$	1	$(Q.y, Q.z)$	4	$(Q.x, Q.y)$
2	$Q.x$	2	$(Q.x, Q.z)$	5	$(Q.x, Q.y)$
3	$Q.x$				

Table 7.2.: `t8_element_boundary_face` for quadrilaterals and hexahedra. Left: For a quadrilateral  $Q$  with anchor node  $(Q.x, Q.y)$  and a face  $f$ , the corresponding anchor node coordinate  $F.x$  of the face line element. Right: For a hexahedron  $Q$  with anchor node  $(Q.x, Q.y, Q.z)$  and a face  $f$ , the corresponding anchor node coordinates  $(F.x, F.y)$  of the face quadrilateral element. In either case, computing the coordinates is equivalent to a projection.

**Remark 7.10.** For the simplicial and hexahedral Morton SFC with maximum refinement level  $\mathcal{L}$ , the anchor node coordinates of an element of level  $\ell$  are integer multiples of  $2^{\mathcal{L}-\ell}$ . Suppose, the maximum level of hexahedral elements is  $\mathcal{L}_1$  and the maximum level of a face boundary quadrilateral element is  $\mathcal{L}_2 \geq \mathcal{L}_1$ , then we will have to multiply a hexahedral coordinate with  $2^{\mathcal{L}_2-\mathcal{L}_1}$  to transform it into a quadrilateral coordinate. We will assume in this section without loss of generality that all element types have the same maximum possible refinement level. Hence, we omit the scaling factor from our equations.

### Quadrilaterals and hexahedra

Let a quadrilateral  $Q$  with anchor node  $(Q.x, Q.y)$  and level  $\ell$  be given. Let furthermore a face index  $f$  be given such that the face is a subface of the root tree. The face element is a line  $F$  with level  $\ell$ . The Computation of its anchor node  $(F.x)$  is a mere coordinate projection, depending on  $f$ . If  $f$  is 0 or 1, then  $F.x = Q.y$ , otherwise  $f$  is 2 or 3 and we obtain  $F.x = Q.x$ .

In 3D we have similar projections, which we list in Table 7.2.

### Triangles and tetrahedra

Since we only construct these face elements for tree boundary faces, we do not need to consider all combinations of element and face number, but only those that occur on the tree boundary. In particular, all possible faces are subfaces of the faces of the root simplex  $S_0$ .

Triangles of type 1 never lie on the root tree boundary, hence we only need to consider type 0 triangles. Let  $T$  be a type 0 triangle with anchor node coordinates  $(T.x, T.y)$  and level  $\ell$ . Furthermore, let a face index  $f$  of a subface of the root tree be given. The corresponding face element  $F$  is a line with level  $\ell$  and anchor node coordinate  $F.x$ . We compute this  $F.x$  from  $T$ 's anchor node as follows. If  $f = 0$ ,



Triangle			Tetrahedron				
type( $T$ )	$f$	$F.x$	type( $T$ )	$f$	Cat.	type( $F$ )	$(F.x, F.y)$
0	0	$T.y$	0	0	1	0	$(T.z, T.y)$
	1	$T.x$		1	1	0	$(T.z, T.y)$
	2	$T.x$		2	2	0	$(T.x, T.z)$
		3		2	0	$(T.x, T.z)$	
			1	0	1	1	$(T.z, T.y)$
			2	2	1	1	$(T.z, T.y)$
			3	-	-	-	—
			4	1	2	1	$(T.x, T.z)$
			5	3	2	1	$(T.x, T.z)$

Table 7.3.: `t8_element_boundary_face` ( $T$ ,  $f$ ) for triangles and tetrahedra. Left: The  $x$  coordinate of the anchor node of the boundary line  $F$  at face  $f$  of a triangle  $T$  in terms of  $T$ 's coordinates. Right: The category and the type of the boundary triangle  $F$  at a face  $f$  of tetrahedron  $T$  as well as the anchor node coordinates  $(F.x, F.y)$ .

then  $F.x = T.y$ . If  $f = 1$ , then  $F.x = T.x$ . Finally, if  $f = 2$  then  $F.x = T.x = T.y$ . We show these cases in Table 7.3.

A tetrahedron that lies on the root tree boundary has type different from 3. In order to compute the boundary face, we further distinguish two categories of such tetrahedra.

**Definition 7.11.** For a tetrahedron  $T$  with  $\text{type}(T) \neq 3$ , we identify two **categories** of root face indices  $g$  regarding how the anchor node  $(F.x, F.y)$  of the corresponding face triangle  $F$  is computed from  $T$ 's anchor node  $(T.x, T.y, T.z)$ :

Faces  $g = 0$  and  $g = 1$  of  $S_0$  lie in the  $(x = 0)$ -plane or the  $(x = z)$ -plane of the coordinate system and form category 1. In this case  $(F.x, F.y) = (T.z, T.y)$ .

Category 2 are the faces  $g = 2$  and  $g = 3$  of  $S_0$ . These lie in the  $(y = 0)$ -plane or the  $(y = z)$ -plane and the anchor node of  $F$  is given by  $(F.x, F.y) = (T.x, T.z)$ .

Depending on the type of a parent tetrahedron its faces fall in one of these categories and have a distinguished triangle type 0 or 1 as subface of  $g \in \{0, 1, 2, 3\}$ .

We display the various categories of faces depending on the type of a simplex in Table 7.3. We also list the anchor node coordinates of the face elements in this table.

### 7.1.3. (iii) Constructing $F'$ from $F$

If we know the tree face index  $g$ , we can look up the corresponding face index  $g'$  of the face in  $K'$  from the coarse mesh connectivity; see Section 6.3.

In order to transform the coordinates of  $F$  to obtain  $F'$  we need to understand how the vertices of the face  $g$  connect to the vertices of the face  $g'$ . Each face's vertices

form a subset of the vertices of the trees as in Figure 6.3. Let  $\{v_0, \dots, v_{n-1}\}$  and  $\{v'_0, \dots, v'_{n-1}\}$  be these vertices for  $g$  and  $g'$  in ascending order, thus  $v_i < v_{i+1}$  and  $v'_i < v'_{i+1}$ . The face-to-face connection of the two trees determines a permutation  $\sigma \in S_n$  such that vertex  $v_i$  connects to vertex  $v'_{\sigma(i)}$ . In theory, there are  $n!$  possible permutations. However, not all of them occur.

**Definition 7.12.** Since we exclude the possibility of trees with negative volume, there is exactly one combination in which the vertices  $v_0$  and  $v'_0$  are connected (c.f. Section 6.1.3). We call the corresponding permutation  $\sigma_0$ .

All the other possible permutations result from rotating the face  $g'$  in this position. This rotation is encoded in the orientation information of the coarse mesh; see Definition 6.2.

Thus, in order to understand the permutation  $\sigma$ , it suffices to know the initial permutation  $\sigma_0$  and the orientation.  $\sigma_0$  is determined by the types of  $K$  and  $K'$  and the face indices  $g$  and  $g'$ . In fact, since the orientation encodes the possible rotations, the only data we need to know is the sign of  $\sigma_0$ .

**Definition 7.13.** Let  $K$  and  $K'$  be two trees of types  $t$  and  $t'$ , and let  $g, g'$  faces of  $K$  and  $K'$  of the same element type. We define the **sign** of  $g$  and  $g'$  as the sign of the permutation  $\sigma_0$ .

$$\text{sign}_{t,t'}(g, g') := \text{sign}(\sigma_0). \quad (7.2)$$

**Remark 7.14.** Note that this definition does not depend on the order of the faces  $g$  and  $g'$ . Since if we switch their roles, the permutation changes to its inverse  $\sigma_0^{-1}$ , thus

$$\text{sign}_{t',t}(g', g) = \text{sign}(\sigma_0^{-1}) = \text{sign}(\sigma_0) = \text{sign}_{t,t'}(g, g'). \quad (7.3)$$

**Remark 7.15.** For hexahedra we can compute the sign of two faces via the tables  $\mathcal{R}, \mathcal{Q}, \mathcal{P}$  from [34, Table 3] as

$$\text{sign}_{\text{hex,hex}}(g, g') = \text{sign}(i \mapsto \mathcal{P}(\mathcal{Q}(\mathcal{R}(g, g'), 0), i)) = -\mathcal{R}(g, g'). \quad (7.4)$$

The permutation in the middle is exactly the permutation  $\sigma_0$ . The argument 0 of  $\mathcal{Q}$  is the orientation of a face-to-face connection, but the result is independent of it, and we could have chosen any other value.

We display the sign for two example tree-to-tree connections, namely tetrahedron to tetrahedron and hexahedron to prism, in Table 7.4.

Using the orientation, the sign, and the face index  $g'$ , we transform the coordinates of  $F$  to obtain the corresponding face  $F'$  as a subface of the face  $G'$  of  $K'$ . For this task we introduce the low-level function

```
t8_element_transform_face (face_element F, orientation o, sign s)
Return the transformed face element F'.
```

$K$ and $K'$ tetrahedra					$K$ hexahedron, $K'$ prism								
		$g$						$g$					
		0	1	2	3			0	1	2	3	4	5
$g'$	0	-1	1	-1	1	$g'$	0	1	-1	-1	1	1	-1
	1	1	-1	1	-1		1	-1	1	1	-1	-1	1
	2	-1	1	-1	1		2	1	-1	-1	1	1	-1
	3	1	-1	1	-1								

Table 7.4.:  $\text{sign}_{t,t'}(g, g')$  from Definition 7.13 for two possible tree-to-tree connections. We obtain these values from Figure 6.3. We do not show the remaining three cases of  $K$  and  $K'$  both being hexahedra or prisms, and  $K$  being a tetrahedron and  $K'$  a prism, which can be obtained similarly.

Line		Quadrilateral			
$o$	$(F'.x)$	$o$	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$	$o$	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$
0	$(F.x)$	0	$\begin{pmatrix} F.x \\ F.y \end{pmatrix}$	2	$\begin{pmatrix} F.y \\ 2^{\mathcal{L}} - F.x - h \end{pmatrix}$
1	$(2^{\mathcal{L}} - F.x - h)$			3	$\begin{pmatrix} 2^{\mathcal{L}} - F.x - h \\ 2^{\mathcal{L}} - F.y - h \end{pmatrix}$
		1	$\begin{pmatrix} 2^{\mathcal{L}} - F.y - h \\ F.x \end{pmatrix}$		

Triangle					
$\text{type}(F)$	$o$	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$	$\text{type}(F)$	$o$	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$
0	0	$\begin{pmatrix} F.x \\ F.y \end{pmatrix}$	1	0	$\begin{pmatrix} F.x \\ F.y \end{pmatrix}$
	1	$\begin{pmatrix} 2^{\mathcal{L}} - F.y - h \\ F.x - F.y \end{pmatrix}$		1	$\begin{pmatrix} 2^{\mathcal{L}} - F.y - h \\ F.x - F.y - h \end{pmatrix}$
	2	$\begin{pmatrix} 2^{\mathcal{L}} - F.x + F.y - h \\ 2^{\mathcal{L}} - F.x - h \end{pmatrix}$		2	$\begin{pmatrix} 2^{\mathcal{L}} - F.x + F.y \\ 2^{\mathcal{L}} - F.x - h \end{pmatrix}$

Table 7.5.: `t8_transform_face` ( $F$ ,  $o$ ,  $s = 1$ ) for lines (top left), quadrilaterals (top right) and triangles (bottom) with sign 1. For values with  $s = -1$  see Table 7.6 and Remark 7.16.

Triangle		Quadrilateral
type( $F$ )	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$	$\begin{pmatrix} F'.x \\ F'.y \end{pmatrix}$
0	$\begin{pmatrix} F.x \\ F.x - F.y \end{pmatrix}$	$\begin{pmatrix} F.y \\ F.x \end{pmatrix}$
1	$\begin{pmatrix} F.x \\ F.x - F.y - h \end{pmatrix}$	

Table 7.6.: `t8_transform_face` ( $F$ ,  $o = 0$ ,  $s = -1$ ) for triangles (left) and quadrilaterals (right). We compute any arbitrary combination of values for  $o$  with  $s = -1$  by first applying `t8_transform_face` ( $F$ ,  $0$ ,  $-1$ ) and then `t8_transform_face` ( $F$ ,  $o$ ,  $1$ ) from Table 7.5.

**Remark 7.16.** The transformation  $o = i, s = -1$  is the same as first  $o = 0, s = -1$  followed by  $o = i, s = 1$ . We thus only need to compute the cases with  $s = 1$  and the specific case  $o = 0, s = -1$ .

Because of  $\sigma_0(0) = 0$ , the sign is always 1 for vertex faces and line faces. Hence, the sign argument is not relevant for the computation of 1D and 2D tree connections.

For the classical and tetrahedral Morton indices we need to compute the anchor node of  $F'$  from the anchor node of the input face  $F$ . We show the computation  $o = i, s = 1$  cases for lines, triangles and quadrilaterals in Table 7.5. Since we transform faces, it is not necessary to discuss the routine for 3 dimensional face element types. We describe the formulas for  $o = 0, s = -1$  for triangles and quadrilaterals in Table 7.6. As we mention in Remark 7.16, we can compute all  $o$  and  $s$  combinations from these two tables. Note that for quadrilaterals and hexahedra `t8_element_transform_face` is equivalent to the internal coordinate transformation in `p4est_transform_face` of the `p4est` library due to (7.4).

#### 7.1.4. (iv) Constructing $E'$ from $F'$

We now have  $E, F, F', K$  and  $K'$  and can construct the neighbor element  $E'$ . For this we use the function

```
t8_element_extrude_face (Face_element F', Tree K', face index g')
Return the element  $E'$  that is a descendant of  $K'$  and has the face  $F'$  at the tree face  $g'$ .
```

`t8_element_extrude_face` has as input a face element and a root tree face index and as output the element within the root tree that has as a boundary face the given face element. How to compute the element from this data depends on the element type and the root tree face. For quadrilaterals, triangles, hexahedra, and tetrahedra with the (TM-)Morton index we show the formulas to compute the anchor node

coordinates of  $E'$  in Table 7.7.

### 7.1.5. A note on vertex/edge-neighbors

Despite the restriction to face-neighbors in this thesis, we are certain that for tree-to-tree neighbors the same method of constructing the lower dimensional element, transforming it into its neighbor and then extruding it to the neighbor element can be applied to vertex and edge neighbors as well. The challenge with these neighbors compared to face-neighbors is that at a single vertex/edge an arbitrary number of trees can be connected. Identifying the correct neighbor trees is, however, a task that entirely relies on the coarse mesh connectivity. Once this is accomplished, the neighbor elements can be constructed by using the techniques described in this section.

## 7.2. Half-size face-neighbors

In order to implement the ghost algorithm for balanced forests from [34], we need to compute half face-neighbors of an element. That is, given an element  $E$  and a face  $f$ , construct the neighbors<sup>2</sup> of  $E$  across  $f$  of refinement level  $\ell(E) + 1$ .

We construct the half face-neighbors in three steps:

- (i) Construct the children  $C_f$  of  $E$  that have a face in  $f$ .
- (ii) For each child  $C_f[i]$  compute the face index  $f_i$  of the face that is a child of  $f$  and a face of  $C_f[i]$ .
- (iii) For each child  $C_f[i]$  compute its face-neighbor across  $f_i$ .

(i) and (ii) are performed by low-level algorithms:

```
t8_element_children_at_face (Element E, face_index f)
```

Returns an array of children of  $E$  that share a face with  $f$ .

```
t8_element_child_face (Element E, child_index i, face_index f)
```

Given an element  $E$ , a child index  $i$ , and a face index  $f$  of a face of  $E$ , compute the index of the  $i$ -th child's face that is a subface of  $f$ .

A typical implementation of `t8_element_children_at_face` would look up the child indices of these children in a table and then construct the children with these indices. The child indices can be obtained from the refinement pattern. For the quadrilateral Morton index, for example, the child indices at face  $f = 0$  are 0 and 2. For a hexahedron the child indices at face  $f = 3$  are 2, 3, 6, and 7. For the TM

---

<sup>2</sup>These do not need to be of half the size. If for example refinement is 1:9 Peano refinement, the neighbors are one third the size.

Quadrilateral from line $F'$				Triangle from line $F'$			
$g'$	$\begin{pmatrix} E'.x \\ E'.y \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \end{pmatrix}$
0	$\begin{pmatrix} 0 \\ F'.x \end{pmatrix}$	2	$\begin{pmatrix} F'.x \\ 0 \end{pmatrix}$	0	$\begin{pmatrix} 2^{\mathcal{L}} - h \\ F'.x \end{pmatrix}$	2	$\begin{pmatrix} F'.x \\ 0 \end{pmatrix}$
1	$\begin{pmatrix} 2^{\mathcal{L}} - h \\ F'.x \end{pmatrix}$	3	$\begin{pmatrix} F'.x \\ 2^{\mathcal{L}} - h \end{pmatrix}$	1	$\begin{pmatrix} F'.x \\ F'.x \end{pmatrix}$		

Hexahedron from quadrilateral $F'$					
$g'$	$\begin{pmatrix} E'.x \\ E'.y \\ E'.z \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \\ E'.z \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \\ E'.z \end{pmatrix}$
0	$\begin{pmatrix} 0 \\ F'.x \\ F'.y \end{pmatrix}$	2	$\begin{pmatrix} F'.x \\ 0 \\ F'.y \end{pmatrix}$	4	$\begin{pmatrix} F'.x \\ F'.y \\ 0 \end{pmatrix}$
1	$\begin{pmatrix} 2^{\mathcal{L}} - h \\ F'.x \\ F'.y \end{pmatrix}$	3	$\begin{pmatrix} F'.x \\ 2^{\mathcal{L}} - h \\ F'.y \end{pmatrix}$	5	$\begin{pmatrix} F'.x \\ F'.y \\ 2^{\mathcal{L}} - h \end{pmatrix}$

Tetrahedron from triangle $F'$ Coordinates			Tetrahedron from triangle $F'$ Type			
$g'$	$\begin{pmatrix} E'.x \\ E'.y \\ E'.z \end{pmatrix}$	$g'$	$\begin{pmatrix} E'.x \\ E'.y \\ E'.z \end{pmatrix}$	$g'$	type( $F'$ )	type( $E'$ )
0	$\begin{pmatrix} 2^{\mathcal{L}} - h \\ F'.y \\ F'.x \end{pmatrix}$	2	$\begin{pmatrix} F'.x \\ 0 \\ F'.y \end{pmatrix}$	0	0	0
				1	1	1
1	$\begin{pmatrix} F'.x \\ F'.y \\ F'.x \end{pmatrix}$	3	$\begin{pmatrix} F'.x \\ 0 \\ F'.y \end{pmatrix}$	1	0	0
				2	0	0
				3	0	0

Table 7.7.: The computation of  $E' = \text{t8\_element\_extrude\_face}(F', T', g')$  for  $T'$  a quadrilateral (top left), triangle (top right), hexahedron (middle), or tetrahedron (bottom). Depending on the anchor node coordinates of  $F'$  and the tree face index  $g'$ , we determine the anchor node of the extruded element  $E'$ . For tetrahedra, we additionally need to compute the type of  $E'$  from  $g'$  and the type of the triangle  $F'$  (bottom right). In the case of a triangle, the type of  $E'$  is always 0, since type 1 triangles cannot lie on a tree boundary.  $h$  refers to the length of the element  $E'$  (resp.  $F'$ ) and is computed as  $2^{\mathcal{L}-\ell}$  where  $\ell$  is the refinement level of  $E'$  and  $F'$ .

Triangle				Tetrahedron				
	$f$				$f$			
type( $T$ )	0	1	2	type( $T$ )	0	1	2	3
0	1,3	0,3	0,1	0	1, 4, 5, 7	0, 4, 6, 7	0, 1, 2, 7	0, 1, 3, 4
1	2,3	0,3	0,2	1	1, 4, 5, 7	0, 5, 6, 7	0, 1, 3, 7	0, 1, 2, 5
				2	3, 4, 5, 7	0, 4, 6, 7	0, 1, 3, 7	0, 2, 3, 4
				3	1, 5, 6, 7	0, 4, 6, 7	0, 1, 3, 7	0, 1, 2, 6
				4	3, 5, 6, 7	0, 4, 5, 7	0, 1, 3, 7	0, 2, 3, 5
				5	3, 5, 6, 7	0, 4, 6, 7	0, 2, 3, 7	0, 1, 3, 6

Table 7.8.: The child indices of all children of an element touching a given face. These indices are needed for `t8_element_children_at_face`. Left: The indices for a triangle  $T$  in dependence on its type and the face index  $f$ . Right: The same data for a tetrahedron  $T$ .

index these indices additionally depend on the type of the simplex. We list all cases in Table 7.8.

The low-level algorithm `t8_element_child_face` can also be described via lookup tables. Its input is a parent element  $E$ , a face index  $f$  and a child index  $i$ , such that the child  $E_i$  of  $E$  has a subface of the face  $f$ . In other words,  $E_i$  is part of the output of `t8_element_children_at_face`. The return value of `t8_element_child_face` is the face index  $f_i$  of the face of  $E[i]$  that is the subface of  $f$ .

For the classical Morton index, the algorithm is the identity on  $f$ , since the faces of child quadrilaterals/hexahedra are labeled in the same manner as those of the parent element. For the TM index for triangles, the algorithm is also the identity, since only triangle children of the same type as the parent can touch a face of the parent and for same type triangles the faces are labeled in the same manner.

For tetrahedra, the algorithm is the identity on those children that have the same type as the parent. However, for each face  $f$  of a tetrahedron  $T$ , there exists a child of  $T$  that has the middle face child of  $f$  as a face. This child has not the same type as  $T$ . For this child the corresponding face value is computed as 0 if  $f = 0$ , 2 if  $f = 1$ , 1 if  $f = 2$ , or 3 if  $f = 3$ .

### 7.3. Finding owner processes of elements

For the ghost algorithm, after we have successfully constructed an element's (half) face-neighbor, we need to identify the owner process of this neighbor.

**Definition 7.17.** Let  $E$  be an element in a (partitioned) forest. A process  $p$  is an **owner** of  $E$  if there exists a leaf  $L$  in the forest such that

1.  $L$  is in the partition of  $p$ , and
2.  $L$  is an ancestor or a descendant of  $E$ .

Note that the owner of an arbitrary element is not unique. Unique ownership is, however, guaranteed for leaf elements and their descendants. Also, each element has at least one owner.

In this section, we describe how to find all owner processes of a given element and how to find those processes that own leaf elements sharing a given face with an element.

### 7.3.1. `t8_forest_owner`

We begin with the algorithm `t8_forest_owner` that determines all owner processes of a given forest element.

**Definition 7.18.** The **first/last descendant** of an element  $E$  is the descendant of  $E$  of maximum refinement level with smallest/largest SFC index.

Since first/last descendants cannot be refined further, they are either a leaf or descendants of a leaf. Hence, they have a unique owner process. See also Figure 7.4 for an illustration. We denote these owners by  $p_{\text{first}}(E)$  and  $p_{\text{last}}(E)$ . Since a forest is always partitioned along the SFC in ascending order, it must hold for each owner process  $p$  of  $E$  that

$$p_{\text{first}}(E) \leq p \leq p_{\text{last}}(E). \quad (7.5)$$

On the other hand, if a process  $p$  fulfills inequality 7.5 and its partition is not empty, then it must be an owner of  $E$ . Furthermore, we conclude that an element has a unique owner if and only if  $p_{\text{first}}(E) = p_{\text{last}}(E)$ .

Each process can compute the SFC index of the first descendant of its first local element. From these SFC indices we build an array of size  $P$ , which is the same on each process. We can then determine the owner process of a first or last descendant by performing a binary search in this array if we combine it with the array of tree offsets. This is the same approach as in [34].

Hence, we can compute all owner processes of an element by constructing its first and last descendant and computing their owners. When we know that an element has a unique owner—for example when it is a leaf element—it suffices to construct its first descendant and compute its owner.

### 7.3.2. Owners at a face

For the `Ghost_v2` algorithm that works on an unbalanced forest—as described in [70] for cubical elements—we will have to identify all owners of leaves at a face of a face-neighbor element of a given element. In contrast to the algorithm `find_range_boundaries` that the authors of [70] use, we introduce the algorithm `t8_owners_at_face`. Given an element  $E$  and a face  $f$ , `t8_owners_at_face` determines the set  $P_E$  of all processes that have leaf elements that are descendants of  $E$  and share a face with  $f$ . It is a recursive algorithm that we now describe in detail.








Figure 7.5.: An example for Algorithm 7.3.1, `t8_forest_owners_at_face`. Left: A triangle element  $E$  with the TM-index as SFC whose descendants are owned by three different processes: 0 (red), 1 (blue), and 2 (green). The owners at the faces are  $\{0, 2\}$  at face 0,  $\{0, 1, 2\}$  at face 1, and  $\{0\}$  at face 2. Right: The iterations of `t8_forest_owners_at_face` at face  $f = 1$ . At first the first and last descendant of  $E$  at  $f$  are constructed. We compute their owner processes 0 and 2, and since their difference is greater one, we continue the recursion. In the second iteration the algorithm is called once for the lower left child and once for the upper child of  $E$ . We determine their first and last descendants at the respective subspace of  $f$ . For the lower left child, the recursion stops since both face descendants are owned by process 0. For the upper child the owner processes are 1 and 2 and since there are no other possible owner processes in between, we stop the recursion as well.

---

**Algorithm 7.3.1:** `t8_owners_at_face` (Forest  $\mathcal{F}$ , element  $E$ , face\_index  $f$ )

---

**Result:** The set  $P_E$  of all processes that own leaf elements that are descendants of  $E$  and have a face that is a subface of  $f$ .

```

1  $P_E \leftarrow \emptyset$ 
2  $fd \leftarrow \text{t8\_element\_first\_desc\_face}(E, f)$  /* First and last descendant of  $E$  at  $f$  */
3  $ld \leftarrow \text{t8\_element\_last\_desc\_face}(E, f)$ 
4  $p_{\text{first}} \leftarrow \text{t8\_forest\_owner}(\mathcal{F}, fd)$  /* The owners of  $fd$  and  $ld$  */
5  $p_{\text{last}} \leftarrow \text{t8\_forest\_owner}(\mathcal{F}, ld)$ 
6 if  $p_{\text{first}} \in \{p_{\text{last}}, p_{\text{last}} - 1\}$  then /* Only  $p_{\text{first}}$  and  $p_{\text{last}}$  are owners of leaves at  $f$  */
7 |   return  $\{p_{\text{first}}, p_{\text{last}}\}$ 
8 else /* There may be other owners. Enter the recursion. */
9 |    $C_f[] \leftarrow \text{t8\_element\_children\_at\_face}(E, f)$ 
10 |   for  $0 \leq i < \text{t8\_element\_num\_face\_children}(E, f)$  do
11 |      $j \leftarrow \text{child\_index}(C_f[i])$  /* The child number relative to  $E$ . */
12 |      $f' \leftarrow \text{t8\_element\_child\_face}(E, j, f)$  /* The face number of the child */
13 |      $P_E \leftarrow P_E \cup \text{t8\_owners\_at\_face}(\mathcal{F}, C_f[i], f')$  /* Recursion */
14 |   return  $P_E$ 

```

---

### Notes on the implementation

In our implementation of `t8_owners_at_face` we take into account that the first and last owners  $p_f$  and  $p_l$  at the current recursion step form lower and upper bounds for the first and last owners in any upcoming recursion step. Thus, we restrict the binary searches in `t8_forest_owner` to the interval  $[p_f, p_l]$  instead of  $[0, P - 1]$ .

We also exploit that the first descendant of an element  $E$  at a face  $f$  is also the first face descendant of  $E$ 's first child at  $f$ . The same holds for the last descendant and the last child at  $f$ . Thus, we reuse the first/last face descendants and owners of  $E$  when we enter the recursion with the first/last child at  $f$ .

## 7.4. The ghost algorithms

For `Ghost_v1` we assume that the forest is balanced and hence we know that all face-neighbor leaf elements of  $E$  have a refinement level between  $\ell(E) - 1$  and  $\ell(E) + 1$ . Therefore, all neighbor elements of  $E$  with level  $\ell(E) + 1$  must have a unique owner process. Thus, to identify the remote processes of  $E$  at  $F$  it suffices to construct the face-neighbors of  $E$  across  $F$  of level  $\ell(E) + 1$  and determine their owner processes. We construct these face-neighbors via the function `t8_forest_half_face_neighbors`; see Section 7.2. We present the complete `Ghost_v1` algorithm in Algorithm 7.4.1.

For `Ghost_v2` we drop the assumption of a balanced forest and thus there is no a priori knowledge about the face-neighbor leaves of  $E$ . To compute  $E$ 's remote processes across  $F$ , we first construct the corresponding face-neighbor  $E'$  of the same-level as  $E$ . For this element we know that it is either a descendant of a forest

leaf, which then has a unique owner, or an ancestor of multiple forest leaves, which could all have different owners. We need to compute only the owners of those descendant/ancestor forest leaves of  $E'$  that touch the face  $F$ . We achieve this with Algorithm 7.3.1 `t8_forest_owners_at_face` that we describe in Section 7.3.2. We show the complete `Ghost_v2` Algorithm in 7.4.2. It uses the function `dual_face` which, given an element  $E$  and a face index  $f$ , simply returns the face index  $f'$  from the neighboring element.

---

**Algorithm 7.4.1: Ghost\_v1 (Forest  $\mathcal{F}$ ) (for balanced forests only)**

---

```

Result: The ghost layer of  $\mathcal{F}$  is constructed.
1 Ghost_init ()
2 for  $K \in \mathcal{F}.trees$  do
3   for  $E \in K.elements$  do
4     for  $0 \leq f < t8\_element\_num\_faces(E)$  do
5        $E'[] \leftarrow t8\_forest\_half\_face\_neighbors(\mathcal{F}, E, f)$ 
6       for  $0 \leq i < t8\_element\_num\_face\_children(E, f)$  do
7          $q \leftarrow t8\_forest\_owner(\mathcal{F}, E'[i])$ 
8         if  $q \neq p$  then
9            $R_p^q = R_p^q \cup \{E\}$ 
10 Ghost_communicate ()
    /* We outsource the init and communication routine, for later reuse */
1 Function Ghost_init
2 for  $0 \leq q < P$  do
3    $R_p^q = \emptyset$ 
4 Function Ghost_communicate
5  $\mathcal{F}.ghosts \leftarrow \emptyset$ 
6 for  $\{q \mid R_p^q \neq \emptyset\}$  do
7   Send  $R_p^q$  to  $q$ 
8   Receive  $R_q^p$  from  $q$ 
9    $\mathcal{F}.ghosts \leftarrow \mathcal{F}.ghosts \cup R_q^p$ 

```

---

## 7.5. Optimizing the runtime of Ghost

The `Ghost_v1` and `Ghost_v2` algorithms that we present here both iterate over all local leaf elements to identify the boundary leaves on the process's boundary. For each leaf we generate all (half) face-neighbors and compute their owners. However, for most meshes only a portion of the leaf elements actually are boundary elements, depending on the surface-to-volume ratio of the process's partition. Since the surface of a volume grows quadratically while the volume itself grows cubically, the number of boundary leaves can become arbitrarily small in comparison to the number of all leaves.

We thus aim to improve the runtime of the algorithms by excluding inner leaves from the iteration. In `p4est` the inner leaves are excluded from the iteration by

---

**Algorithm 7.4.2: Ghost\_v2 (Forest  $\mathcal{F}$ )**

---

**Result:** The ghost layer of  $\mathcal{F}$  is constructed.

```
1 Ghost_init ()
2 for  $K \in \mathcal{F}.trees$  do
3   for  $E \in T.elements$  do
4     for  $0 \leq f < t8\_element\_num\_faces(E)$  do
5        $E' \leftarrow t8\_forest\_face\_neighbor(\mathcal{F}, E, f)$ 
6        $f' \leftarrow dual\_face(E, E', f)$ 
7        $P_{E'} \leftarrow t8\_forest\_owners\_at\_face(\mathcal{F}, E', f')$ 
8       for  $q \in P_{E'}$  do
9         if  $q \neq p$  then
10           $R_p^q = R_p^q \cup \{E\}$ 
11 Ghost_communicate ()
```

---

checking for each quadrilateral/hexahedron whether its  $3 \times 3$  neighborhood, thus all same-level face-(edge-/vertex-)neighbors, are process local. Since this approach particularly uses geometrical properties of the quadrilaterals/hexahedra and of the Morton SFC, it is not practical for our element-type independent approach.

To exclude the inner leafs, we exchange the leaf iteration with a top-town search using the recursive approach from [70]. Starting with a tree's root element, we check whether it may have boundary leaf descendants, and if so, we create the children of the element and continue recursively. If we reach a leaf element, we check whether it is a boundary element—and if so for which processes—in the way described in the previous section. This approach allows us to terminate the recursion as soon as we reach an element that lies completely within the process's partition, thus saving the iteration over all descendant leaves of that element.

We now discuss the details of the top-down search and how we use it to improve the ghost algorithm.

### 7.5.1. The recursive top-down search

In [70] the authors present the general recursive `search` algorithm for octree AMR, which easily extends to arbitrary tree-based AMR. The setting is that we search a leaf or a set of leaves in a forest that satisfy given conditions. One numerical example for such a search arises in semi-Lagrangian advection solvers [3, 99]. To interpolate the values of an advected function  $\phi_t$  at time  $t$ , each grid point  $x_i$  is tracked back in time to its previous position  $\hat{x}_i$  at  $t - \Delta t$ . This point  $\hat{x}_i$  lies in a leaf element  $E_i$  of the forest and an element-local Hermite interpolation with the values of  $\phi_{t-1}$  is used to determine the value  $\phi_t(x_i)$ . Thus, in each time step, we have to search the forest for the leaf elements  $\{E_i\}$  given the points  $\{\hat{x}_i\}$ .

In our case, we apply `search` to the problem of identifying all leaf elements at a process's boundary. The `search` algorithm has been shown to be especially efficient when looking for multiple matching leaves at once [70], which is the case in our

setting.

As presented in [70] the idea of search is to perform a recursive top-down traversal for each tree by starting with the root element of that tree and recursively creating its children until we end up with a leaf element. On each intermediate element we call a user-provided callback function which returns true only if the search should continue with this element. If the callback returns false, the recursion for this element stops and its children are excluded from the search. If the search has reached a leaf element, the callback also performs the desired operations if the leaf matches the search.

For our ghost algorithm the callback returns false for elements that lie entirely within the process's domain, thus excluding possibly large areas from the search and hence speeding up the computation. Once a leaf element is reached, we check whether it is a boundary element or not. Thus, we iterate over the leaf's faces and compute the owners at the respective neighbor faces as in the inner for loop of Algorithm 7.4.2.

We show our version of `search` in Algorithm 7.5.1. It is a simplified version of Algorithm 3.1 in [70] without queries, since we do not need these for `Ghost`. We also use the function `split_array` from [70]. This function takes as input an element  $E$  and an array  $L$  of (process local) leaf elements in  $E$ , sorted in SFC order. `split_array` returns a set of arrays  $\{M[i]\}$ , such that for the  $i$ -th child  $E_i$  of  $E$  the array  $M[i]$  contains exactly the leaves in  $L$  that are also leaves of  $E_i$ . Thus,  $L = \bigcup_i M[i]$ .

For a search of the complete forest, we iterate over all trees and in each tree we compute the finest element  $E$  such that all tree leaves are still descendants of  $E$ . We compute  $E$  as the nearest common ancestor of the first and last leaf element of the tree. With this  $E$  and the leaf elements of the tree, we call `element_search`. See Algorithm 7.5.2.

---

**Algorithm 7.5.1:** `element_search` (Element  $E$ , Leaf elements  $L$ , Callback `Match`) [See Algorithm 3.1 in [70]]

---

```

Result: Match is called with  $E$  as input. If the result is true, we continue recursively
with  $E$ 's children.
1 if  $L = \emptyset$  then
2   return
3  $\text{isLeaf} \leftarrow L = \{E\}$           /* Boolean to determine whether  $E$  is a leaf element */
4 if Match( $E$ ,  $\text{isLeaf}$ ) and not  $\text{isLeaf}$  then          /* Decide whether to continue
recursion */
5    $M[] \leftarrow \text{split\_array}(L, E)$ 
6    $C[] \leftarrow \text{t8\_element\_children}(E)$ 
7   for  $0 \leq i < \text{t8\_element\_num\_children}(E)$  do
8     return element_search( $C[i]$ ,  $M[i]$ , Match)

```

---

---

**Algorithm 7.5.2:** `t8_forest_search` (Forest  $\mathcal{F}$ , Callback Match)

---

**Result:** `element_search` is called on each tree.

```
1 for  $K \in \mathcal{F}.trees$  do
2    $E_1 \leftarrow \text{first\_tree\_element}(\mathcal{F}, K)$            /* First and last local leaf */
3    $E_2 \leftarrow \text{last\_tree\_element}(\mathcal{F}, K)$            /* in the tree */
4    $E \leftarrow \text{t8\_element\_nearest\_common\_ancestor}(E_1, E_2)$ 
5    $L \leftarrow \text{tree.leaves}(\mathcal{F}, K)$                  /* Array of tree leaves */
6   element_search ( $E, L, \text{Match}$ )
```

---

## 7.5.2. The optimized Ghost algorithm

We use `forest_search` for an optimized version of `Ghost`. When iterating over all leaves of the forest and checking the neighbors for each one, a lot of these elements are in the interior domain of the process. By using search we can exclude a set of interior leaves as soon as the search recursion enters an ancestor that is completely in the interior of the domain.

We show our callback algorithm `t8_ghost_match` in Algorithm 7.5.3, which works as follows. If the element  $E$  which is passed to `t8_ghost_match` is not a leaf element, we check whether the element and all of its possible face-neighbors are owned by the current process. For the element's owners, we do not call the function `t8_forest_owner`, but instead save runtime by computing the first and last process that own leaves of the element and checking whether they are equal. For these computations we construct  $E$ 's first and last descendant. Analogously, for the owners at the neighbor faces we compute the first and last owner processes. If for  $E$  the first and last process is  $p$  and at each face-neighbor the first and last owner at the corresponding face is also  $p$ ,  $E$  is an inner element and cannot have any boundary leaves as descendants. Thus, we return 0 and the search does not continue for the descendants of  $E$ .

If  $E$  is a leaf element, then it may or may not be a boundary element. We thus compute all owner processes for all face-neighbors using `t8_forest_owners_at_face` and add  $E$  as a boundary element to all of these that are not  $p$ .

### Implementation details

For each child  $C$  of an element  $E$  the ranks  $p_{\text{first}}(E)$ ,  $p_{\text{last}}(E)$ ,  $p_{\text{first}}(E, f)$ , and  $p_{\text{last}}(E, f)$  serve as lower and upper bounds for the corresponding ranks for  $C$ . Thus, in our implementation of `ghost_match` in `t8code`, we store these ranks for each recursion level reducing the search range for the binary owner search for  $C$  from  $[0, P - 1]$  to  $[p_{\text{first}}(E), p_{\text{last}}(E)]$ , and to  $[p_{\text{first}}(E, f), p_{\text{last}}(E, f)]$  for the faces. To compute these bounds it is necessary to always enter the `for`-loop in Line 14, even though we do not exercise this in Algorithm 7.5.3.

---

**Algorithm 7.5.3: t8\_ghost\_match** (Element  $E$ , Bool isLeaf)

---

**Result:** If  $E$  is a leaf, compute the owners of the face-neighbors and add to the sets  $R_p^g$ .  
If not, query whether all descendants of  $E$  and all face-neighbors are owned by  $p$ .

```
1 if isLeaf then                               /* E is a leaf. Compute the owners at */
2   for  $0 \leq f < \text{t8\_element\_num\_faces}(E)$  do /* the face and add E as boundary.
   */
3      $E' \leftarrow \text{t8\_forest\_face\_neighbor}(\mathcal{F}, E, f)$ 
4      $f' \leftarrow \text{dual\_face}(E, E', f)$ 
5      $P_{E'} \leftarrow \text{t8\_forest\_owners\_at\_face}(\mathcal{F}, E', f')$ 
6     for  $q \in P_{E'}$  do
7       if  $q \neq p$  then
8          $R_p^g = R_p^g \cup \{E\}$ 
9 else                                           /* E is not a leaf. */
10   $p_{\text{first}}(E) \leftarrow \text{t8\_element\_first\_owner}(E)$ 
11   $p_{\text{last}}(E) \leftarrow \text{t8\_element\_last\_owner}(E)$ 
12  if  $p_{\text{first}}(E) > p$  or  $p_{\text{last}}(E) < p$  then /* No leaf of E is owned by p */
13    return 1
14  for  $0 \leq f < \text{t8\_element\_num\_faces}(E)$  do
15     $E' \leftarrow \text{t8\_forest\_face\_neighbor}(\mathcal{F}, E, f)$ 
16     $f' \leftarrow \text{dual\_face}(E, E', f)$ 
17     $p_{\text{first}}(E', f') \leftarrow \text{t8\_first\_owner\_at\_face}(\mathcal{F}, E', f')$ 
18     $p_{\text{last}}(E', f') \leftarrow \text{t8\_last\_owner\_at\_face}(\mathcal{F}, E', f')$ 
19    if  $p_{\text{first}}(E', f') \neq p$  or  $p_{\text{last}}(E', f') \neq p$  then
20      return 1 /* Not all face-neighbor leaves are owned by p */
21  if  $p_{\text{first}}(E) = p_{\text{last}}(E) = p$  then
22    return 0
23 return 1
```

---

---

**Algorithm 7.5.4: Ghost\_v3** (Forest  $\mathcal{F}$ )

---

**Result:** The ghost layer of  $\mathcal{F}$  is constructed.

```
1 Ghost_init ()
2 t8_forest_search( $\mathcal{F}$ , t8_ghost_match)
3 Ghost_communicate ()
```

---



## 7.6. Numerical comparison of the ghost versions

To verify that the additional complexity of implementing the top-down search is worth the effort, we perform runtime tests of the different ghost methods.

We perform tests with hexahedral and tetrahedral elements, each time on a unit cube geometry. For hexahedra the unit cube is modeled with a single tree and for tetrahedra with six trees with a common diagonal as in Figure 4.2. For each element type we run two types of tests, one with a uniform mesh and one with an adaptive mesh, where we use a regular refinement pattern, refining every third element (in SFC order) recursively from level  $\ell$  up to a level  $\ell + k$ ; see Figure 7.6.

Since we are interested in comparing the algorithms and not in their particular extreme scaling behavior, we run the tests on 1024 MPI ranks on JUQUEEN [73]. We refer to Chapter 8 for more elaborate scaling tests of **Ghost** on significantly more ranks (up to 458k). For the tests in this section we use 64 compute nodes with 16 cores and 16 GB memory each. We use 1 rank per core, thus 16 MPI ranks per node. We display our results in Table 7.9 showing runtime results for uniform levels  $\ell$  equal to 9, 8 and 4, and adaptive levels  $\ell$  equal to 8, 7, and 3 (tetrahedra), respectively 4 (hexahedra), with  $k = 2$ .

As expected, the iterative versions of **Ghost** scale linearly with the number of elements. The improved version of **Ghost**, however, scales with the number of ghost elements, which grows less quickly compared to the number of elements. From this we conclude that we indeed skip most of the elements that do not lie on the boundary of a process's domain. The improved version shows overall a significantly better performance and is up to a factor of 23.7 faster (adaptive tetrahedra, level 8) than the iterative version. For smaller or degraded meshes where the number of ghosts is on the same order as the number of leaf elements, the improved version shows no disadvantage compared to the iterative version. This shows that we do not lose runtime to the **Search** overhead, even if each element is a boundary element. For small meshes all algorithms show negligible runtime on the order of milliseconds.

We conclude that the ghost version with top-down search is the ideal choice among the three version that we discuss. From now on, we use this algorithm for all tests.




Figure 7.6.: We compare the different implementations of **Ghost** by testing them on a unit cube geometry with 1024 MPI ranks. Left: The adaptive mesh with minimum level  $\ell = 3$  for hexahedra. We refine every third element in SFC order and repeat the process once with the refined elements. Right: For an adaptive computation with 4 MPI ranks, we show the local leaf elements of the process with MPI rank 1 (red) and its ghost elements (blue, transparent).

tetrahedra						
	uniform			adaptive		
$\ell$	9	8	4	8-10	7-9	3-5
elements/proc	786,432	98,304	24	1,015,808	126,976	31
ghosts/proc	32,704	8,160	30	31,604	8,137	56
Ghost_v1 [s]	172.3	21.64	7.99e-3	-	-	-
Ghost_v2 [s]	129.6	16.19	5.93e-3	167.94	20.88	8.10e-3
Ghost_v3 [s]	7.41	1.75	5.01e-3	7.08	1.69	8.12e-3

hexahedra						
	uniform			adaptive		
$\ell$	9	8	4	8-10	7-9	4-6
elements/proc	131,072	16,384	4	169,301	21,162	41
ghosts/proc	8,192	2,048	8	7,681	1,913	30
Ghost_v1 [s]	29.51	3.742	2.87e-3	-	-	-
Ghost_v2 [s]	18.25	2.302	2.32e-3	23.79	2.964	8.01e-3
Ghost_v3 [s]	3.14	0.711	2.90e-3	2.81	0.649	8.12e-3

Table 7.9.: Runtime tests for the three different `Ghost` algorithms that we describe in this chapter. We run the tests on JUQUEEN with 1024 MPI ranks and 16 MPI ranks per compute node. The domain geometry is a unit cube modeled by one tree in the hexahedral case and six trees in the tetrahedral case. With each element type we test a uniform level  $\ell$  mesh and a mesh that adapts every third element of a uniform level  $\ell$  mesh up to level  $\ell + 2$ ; see Figure 7.6. The different `Ghost` methods are: `Ghost_v1` that works on balanced forests only; `Ghost_v2` that works also on unbalanced forests; `Ghost_v3` that utilizes `search` to improve the runtime. Since the adaptive forests are not balanced, we do not test `Ghost_v1` in that case. The table at the top shows the mesh sizes and runtimes for tetrahedra while the bottom table shows the data for hexahedra. We observe that our new `Ghost_v3` is superior to the other versions by a factor of up to 23 and scales with the number of ghosts and not with the number of elements.

## 8. 2:1 Balance

It is common for AMR applications, such as finite element and finite volume solvers, to rely on a forest mesh that is 2:1 balanced, meaning that each element only has neighbors whose levels differ by at most one ( $\pm 1$ ) from the element's level [33, 41, 85, 104]; see also Definition 7.4. This restricts the number and configurations of hanging nodes/edges/faces that can occur, simplifying the necessary interpolation schemes and reducing the number of neighboring processes.

However, mesh refinement and coarsening in applications is usually driven by some kind of error estimator and/or geometric constraints and such adaptation rules may not produce balanced meshes on their own. It can become a significant challenge to change the adaptation rule such that it produces balanced meshes that still respect the desired constraints.

We thus aim to decouple the operation of balancing a forest from the adaptation routine. An application can then call **Adapt** to modify the mesh and optionally call the algorithm **Balance** afterwards to reestablish a balanced mesh. In this section we discuss our implementation of **Balance**.

Note that there are also applications that can handle arbitrarily hanging nodes at elements with refinement level difference greater than one [61, 132]. In general it is up to an application whether to use **Balance** or not.

The algorithm **Balance** gets as input a forest that may be unbalanced and modifies it by successively refining leaf elements such that it becomes balanced. **Balance** should not coarsen any elements in order to guarantee that an application can keep its desired accuracy.

As with **Ghost**, we distinguish between corner-balance, edge-balance, and face-balance, regarding the different possible neighbor connections. As we mention before, we restrict ourselves to face-neighbors and thus we consider face-balance here, sometimes also referred to as 1-balance [69].

In **Balance**, a leaf element with a large refinement level that is surrounded by leaves of smaller refinement levels can trigger refinement of leaves over large regions that may stretch across multiple process boundaries; see Figure 8.1. This is one of the reasons why **Balance** was shown to be the most expensive high-level algorithm [34]. The relatively high run time costs of **Balance** have sparked efforts to optimize and speed up the algorithm [69].

In this thesis, we restrict ourselves to a straightforward implementation of **Balance** via the existing algorithms **Adapt** and **Ghost**. The idea is similar to the ripple algorithm from [143, 144]. We see the implementation that we give here as a feasibility study of **Balance** for meshes with arbitrary element types and do not claim to achieve an optimal runtime. We thus also refer to our algorithm as **Ripple-balance**.




Figure 8.1.: An example for face-balance with a triangular forest and 5 MPI ranks. Left: An unbalanced and fairly equally partitioned forest with two trees. Right: After balancing the forest, each leaf is refined in such a way that no two face-neighboring leaf elements have a level difference of more than one. Note that the finest leaf elements reside on rank 2 and influence the refinement of leaves on ranks 0, 3, and 4. We also observe that the load-balance has been disturbed. To prevent this, we may add additional partitioning steps during and/or after `Balance`.

Implementing an optimized algorithm in the spirit of [69] remains a challenge for future work.

## 8.1. Finding leaf descendants of an element

A key ingredient of our primitive—yet functional—version of `Balance` is to identify those leaves that have neighbors of greater refinement level than the leaf’s level plus one. We thus need to query for a leaf element  $E$  in a forest  $\mathcal{F}$  whether there exists a face-neighbor leaf of  $E$  in  $\mathcal{F}$  with larger refinement level than  $\ell(E) + 1$ . To this end, we construct  $E$ ’s face-neighbor elements of level  $\ell(E) + 1$  via the function `t8_element_half_face_neighbors` from Section 7.2. For each of these half face-neighbors  $E'$ , we check whether there exists a true descendant—i.e. a descendant that is not  $E'$  itself—that is a local leaf element or ghost element in  $\mathcal{F}$ .

In order to perform this check, we create the last descendant  $D$  of  $E'$  and search

for an element  $L \in \mathcal{F}.\text{elements} \cup \mathcal{F}.\text{ghosts}$  such that

$$\mathcal{I}(E') < \mathcal{I}(L) \leq \mathcal{I}(D). \quad (8.1)$$

Here,  $\mathcal{I}$  is the forest wide SFC index as in Section 3.4. Such an  $L$ , if found, is a descendant of  $E'$  because of the properties of  $\mathcal{I}$ . Also,  $L$  is not  $E'$  itself since then  $\mathcal{I}(E') = \mathcal{I}(L) \not< \mathcal{I}(L)$ . If no such  $L$  exists, we know that  $E'$  does not have a true descendant in the leaves or ghosts.

We show the algorithm `t8_forest_leaf_desc_exists` in Algorithm 8.1.1. To search for an  $L$  that fulfills (8.1) in Lines 2 and 5 we need to perform a binary search for an element  $D$  in a sorted (in SFC-order) array  $A$  of  $n$  elements, where we cannot guarantee that  $D \in A$ , but expect as result the largest integer  $i$  such that  $\mathcal{I}(A[i]) \leq \mathcal{I}(D)$ . If we assume that  $\mathcal{I}(A[0]) \leq \mathcal{I}(D)$ , this binary search is possible: Choose bounds  $l = 0, h = n - 1$  and a guess  $g = (l + h + 1)/2$ ; if  $\mathcal{I}(A[g]) > \mathcal{I}(D)$  then set  $h = g - 1$ , else  $\mathcal{I}(A[g]) \leq \mathcal{I}(D)$  and we set  $l = g$ . Start again with  $g = (l + h + 1)/2$  and iterate. We show this operation in the function `binary_search` in Algorithm 8.1.1.

If the assumption  $\mathcal{I}(A[0]) \leq \mathcal{I}(D)$  is not fulfilled, an  $L$  that satisfies equation (8.1) does not exist. We can check this in constant time and the search returns  $L = A[0]$ , for which the check  $\mathcal{I}(E') < \mathcal{I}(L) \leq \mathcal{I}(D)$  from Line 3 or Line 6 fails as expected.

---

**Algorithm 8.1.1:** `t8_forest_leaf_desc_exists` (Forest  $\mathcal{F}$ , Element  $E'$ )

---

**Result:** True if a leaf in  $\mathcal{F}.\text{elements}$  or  $\mathcal{F}.\text{ghosts}$  exists that is a true descendant of  $E'$ . False otherwise.

```

1  $D \leftarrow \text{t8.element.last.descendant}(E')$ 
2  $L \leftarrow \text{binary\_search}(\mathcal{F}.\text{elements}, D)$            /* Search in the local elements */
3 if  $\mathcal{I}(E') < \mathcal{I}(L) \leq \mathcal{I}(D)$  then
4   | return True
5  $L \leftarrow \text{binary\_search}(\mathcal{F}.\text{ghosts}, D)$          /* Search in the ghost elements */
6 if  $\mathcal{I}(E') < \mathcal{I}(L) \leq \mathcal{I}(D)$  then
7   | return True
8 return False
```

**Function** `binary_search` (Array  $A$ , Element  $D$ )

**Result:** If  $\mathcal{I}(A[0]) \leq \mathcal{I}(D)$ ,  $A[i]$  for the largest index  $i$  such that  $\mathcal{I}(A[i]) \leq \mathcal{I}(D)$ , otherwise  $A[0]$ .

```

1 if  $\mathcal{I}(A[0]) > \mathcal{I}(D)$  then
2   | return  $A[0]$ 
3  $l \leftarrow 0$ 
4  $h \leftarrow A.\text{length} - 1$ 
5 while  $l < h$  do
6   |  $g \leftarrow \frac{l+h+1}{2}$ 
7     | if  $\mathcal{I}(A[g]) \leq \mathcal{I}(D)$  then  $l = g$ 
8     | else  $h = g - 1$ 
9 return  $A[g]$ 
```

---

## 8.2. The Ripple-balance algorithm

Our ripple version of **Balance** (Algorithm 8.2.1) is an iterative one. In each iteration, we construct a new forest  $\mathcal{F}_{i+1}$  from the current forest  $\mathcal{F}_i$ , starting with the original forest  $\mathcal{F}_0$  that we want to balance. In each iteration, we check for each leaf  $E$  of  $\mathcal{F}_i$  whether there are face-neighbors of  $E$  in the local leaves or ghosts of  $\mathcal{F}_i$  with a larger refinement level than  $\ell(E) + 1$ . If so, we refine the element  $E$  and add the children to the new forest  $\mathcal{F}_{i+1}$ , otherwise, we add the element  $E$  to  $\mathcal{F}_{i+1}$ . For this check we use the function `t8_element_leaf_desc_exists` that we describe in the previous Section. We repeat these refinement steps until the forest mesh on each process does not change anymore.

We use the ghost layer of  $\mathcal{F}_i$  to take into account that a local leaf may need to be refined if a face-neighbor leaf on a neighboring process has a larger refinement level. Thus, in each iteration, we call **Ghost** for the newly constructed forest  $\mathcal{F}_i$ . Note that we need the unbalanced version of **Ghost** here (`Ghost_v2` or `Ghost_v3`).

In Algorithm 8.2.1 we explicitly write down the element loop; however, in the actual implementation we replace it by a call to **Adapt** with the appropriate callback function.

**Proposition 8.1.** *Algorithm 8.2.1 terminates and produces a balanced forest  $\mathcal{F}^*$ .*

*Proof.* The algorithm terminates, if on each process no leaf element is refined any longer. We now show that we eventually reach this status. Let  $m_i$  be the maximum refinement level of all global leaf elements in  $\mathcal{F}_i$  (across all processes). We claim  $m_i = m_0$ . Each leaf  $E$  in  $\mathcal{F}_i$  is either a leaf element of  $\mathcal{F}_{i-1}$  or a child of a leaf element  $\hat{E}$  in  $\mathcal{F}_{i-1}$ . In the first case  $\ell(E) \leq m_{i-1}$  by definition. Let us consider the second case: The element  $E$  of  $\mathcal{F}_i$  is a child of  $\hat{E}$  which is refined in iteration  $i - 1$ . Hence,  $\ell(E) = \ell(\hat{E}) + 1$ . Since  $\hat{E}$  is refined in iteration  $i - 1$ , there exists a face-neighbor  $E'$  of  $\hat{E}$  in  $\mathcal{F}_{i-1}$  with  $\ell(E') > \ell(\hat{E}) + 1 = \ell(E)$ . Since  $m_{i-1} \geq \ell(E')$  we obtain  $m_{i-1} \geq m_i$  and since no leaf is coarsened, we get the equality  $m_{i-1} = m_i$ , proving our claim.

Therefore, the maximum refinement level of  $\mathcal{F}_i$  is bounded by  $m_0$  and because no leaf is coarsened there must exist a final step  $i^*$  in which no element is refined anymore. Hence, the algorithm terminates after  $i^*$  steps.

Let  $\mathcal{F}^* = \mathcal{F}_{i^*}$  be the result of **Ripple-balance**. Since in the last step no element is refined,  $\mathcal{F}^*$  is an exact copy of  $\mathcal{F}_{i^*-1}$ . Thus, for each leaf  $E$  in  $\mathcal{F}^*$  we know that there exists no face-neighbor  $E'$  with level  $\ell(E') > \ell(E) + 1$ , since otherwise  $E'$  would be refined in iteration  $i^*$  and thus the algorithm would continue with  $i^* + 1$ . Suppose that a leaf  $E$  exists in  $\mathcal{F}^*$  with a face-neighbor leaf  $E'$  with level  $\ell(E') < \ell(E) - 1$ , then  $E'$  has a face-neighbor  $E'' = E$  in  $\mathcal{F}^*$  with  $\ell(E'') > \ell(E') + 1$ , which is a contradiction.

Thus,  $\mathcal{F}^*$  fulfills the balance condition.  $\square$

**Remark 8.2.** As we observe in Figure 8.1, repartitioning of the forest may be necessary after **Ripple-balance**. In order to prevent the algorithm to produce largely

---

**Algorithm 8.2.1: Ripple-balance (Forest  $\mathcal{F}_0$ )**

---

**Result:** A new forest  $\mathcal{F}^*$  consisting of (possibly) refined elements of  $\mathcal{F}_0$ , such that  $\mathcal{F}^*$  fulfills the face-balance condition.

```
1 done  $\leftarrow$  0
2  $i \leftarrow$  0
3 while not done do
4     done  $\leftarrow$  1      /* We are done if no element has to be refined any more */
5      $\mathcal{F}_i \leftarrow$  Partition ( $\mathcal{F}_i$ )      /* Partition the forest (optional) */
6     Ghost ( $\mathcal{F}_i$ )      /* Create the ghost layer */
7      $\mathcal{F}_{i+1} \leftarrow \mathcal{F}_i$ 
8     for  $E \in \mathcal{F}_i$ .elements do
9         refine_flag  $\leftarrow$  0
10        for  $0 \leq f < \text{t8\_element\_num\_faces}(E)$  do
11             $E'[] \leftarrow \text{t8\_forest\_half\_face\_neighbors}(\mathcal{F}_i, E, f)$ 
12            for  $0 \leq i < \text{t8\_element\_num\_face\_children}(E, f)$  do
13                if  $\text{t8\_forest\_leaf\_desc\_exists}(\mathcal{F}_i, E'[i])$  then
14                    refine_flag  $\leftarrow$  1      /* Mark  $E$  for refinement */
15                    goto 16      /* No need to check the remaining neighbors */
16            if refine_flag then      /* Refine  $E$  if necessary */
17                 $\mathcal{F}_{i+1}.$ elements  $\leftarrow \mathcal{F}_{i+1}.$ elements  $\setminus \{E\} \cup \text{t8\_element\_children}(E)$ 
18                done  $\leftarrow$  0
19     $i \leftarrow i + 1$ 
20    MPI_Allreduce (done, MPI_LAND) /* Logical 'and' of all values of done */
    /* on the different MPI ranks */
21 return  $\mathcal{F}_i$ 
```

---



imbalanced loads on the different ranks, we may also repartition each intermediate forest before we start the next iteration as we do in Line 5 of the algorithm. The resulting forest  $F^*$  is then partitioned as a consequence.

**Remark 8.3.** Even though we might not refine local elements of a process  $p$  in one iteration, thus keeping the variable `done` set to `true`, changes of the forest in neighboring processes may render it necessary that we need to refine elements on  $p$  in later iterations. For this reason we need to compute the logical 'and' of all `done` values on all processes, hence the `MPI_Allreduce` call in Line 20. Figure 8.1 shows an example for this situation. Here, the elements on process  $p = 4$  (in dark red) do not change in the first two iterations of `Ripple-balance`, but they are refined multiple times eventually.

### 8.3. Numerical results

In this section, we present numerical results for the `Ghost` and `Ripple-balance` routines. All results are obtained with the `t8_time_forest_partition` example of `t8code` version 0.3. We perform the tests on the JUQUEEN supercomputer [73] and use 16 MPI ranks per compute node throughout.

#### 8.3.1. The test case

In the test we use a similar setting to the test in Section 6.4.1 for coarse mesh partitioning. We start with a uniform forest of level  $\ell$  and refine it in a band along an interface defined by a plane to level  $\ell + k$ . We then call `Ripple-balance` to establish a 2:1 balance among the elements and we create a layer of ghost elements with `Ghost` afterwards. The interface moves through the domain in time in direction of the plane's normal vector. In each time step we adapt the mesh, such that we coarsen elements outside of the band to level  $\ell$  and refine within the band to level  $\ell + k$ . We then repeat the `Ripple-balance` and `Ghost` calls. As opposed to the test in Section 6.4.1, we take the unit cube as our coarse mesh geometry. We run the test once with a hexahedral mesh consisting of one tree and once with a tetrahedral mesh of six trees forming a unit cube as in Figure 4.2 in Section 4.1.1.

We choose the normal vector  $\frac{3}{2} (1, 1, \frac{1}{2})^t$ , and we choose  $\frac{1}{4}$  as width of our refinement band. We move the refinement band with speed  $v$  and scale the time step  $\Delta t$  with the refinement level as

$$\Delta t(\ell) = \frac{C}{2^{\ell} v}, \quad (8.2)$$

$C$  being the CFL-number. It is a measurement for the width of the band of level  $\ell$  elements that will be refined to level  $k$  in the next time step. We set  $C = 0.8$  and choose  $v$  such that  $\frac{1}{v} = 0.64$ . We start the band at position  $x_0(\ell) = 0.56 - 2.5\Delta t(\ell)$  and measure up to 5 time steps.

Thus, for level  $\ell$  and band width  $k$ , we use the program call `t8_time_forest_partition -c MESH -n1 -l\ell -rk -x  $x_0(\ell)$  -X  $x_0(\ell) + 0.25$  -C 0.8 -T  $6\Delta t(\ell)$  -gbo`. Here,




Figure 8.2.: We test **Ghost** and **Ripple-balance** on a unit cube geometry consisting of six tetrahedral trees (left) or one hexahedral tree (right). Starting with a uniform level  $\ell$ , we refine the forest in a band around a plane to level  $\ell + k$ . We then balance the forest and create the ghost layer. In the next time step, the band moves in the direction of the plane's normal vector and we repeat the steps, coarsening previously fine forest elements if they now reside outside of the band. We show the forest after **Ripple-balance** at time step  $t = 2\Delta t(\ell)$  for two different configurations. Left: Tetrahedral elements with  $\ell = 3$ ,  $k = 2$ . In total we have 56,566 tetrahedral elements. Right: Hexahedral elements with  $\ell = 4$ ,  $k = 2$ , summing up to 78,100 hexahedral elements in total. The color represents the refinement level. We draw level  $\ell$  elements opaque.

MESH stands for a file storing the coarse mesh, i.e. either the tetrahedralized unit cube or the unit cube of one hexahedron tree. These coarse meshes can also be generated with the `t8_cmesh_new_hypercube` function of `t8code`. The settings `-g` and `-b` tell the program to construct the ghost layer, and to balance the forest mesh after adaptation. The `-o` setting disables output of visualization files. We refer to Figure 8.2 for an illustration of the setting.

### 8.3.2. Strong scaling

We run a strong scaling test with tetrahedral elements and refinement parameters  $\ell = 8$ ,  $k = 2$  on 8,192 up to 131,072 MPI ranks, increasing the process count by a factor of 2 in each step. We list the runtimes at time  $t = 4\Delta t$  for **Ghost** and **Ripple-balance** in Table 8.1, and plot them together with **Partition** in Figure 8.3.

Since the runtime of **Ripple-balance** depends on the number of process-local leaf elements, we expect it to drop by a factor of 2 if we double the number of processes and thus divide the number of elements per process in two. In particular,

Tetrahedral case with $\ell = 8, k = 2, C = 0.8$ at $t = 4\Delta t$						
$P$	$E/P$	$G/P$	Ripple-balance		Ghost	
			Time [s]	Par. Eff.	Time [s]	Par. Eff.
8,192	234,178	17,946	687.0	100.0%	3.25	100.0%
16,384	117,089	11,311	336.2	102.1%	2.12	96.6%
32,768	58,545	7,184	161.2	106.5%	1.27	102.4%
65,536	29,272	4,560	78.3	109.6%	0.79	104.5%
131,072	14,636	2,859	37.7	113.8%	0.52	99.5%

Table 8.1.: The results for strong scaling of **Ripple-balance** and **Ghost** with tetrahedral elements. The problem parameters are  $\ell = 8, k = 2$ , and  $C = 0.8$  with  $\Delta t$  according to (8.2). We show the runtimes of time step  $t = 4\Delta t$ . After **Ripple-balance** the mesh consists of approximately  $1.91e9$  Tetrahedra. In addition to the runtimes, we show the number of elements per process,  $E/P$ , and ghosts per process,  $G/P$ . We also compute the parallel efficiency of **Ripple-balance** and **Ghost** according to (8.3) and (8.4) in reference to the run with 8,192 processes. We observe a more than ideal scaling for **Ripple-balance** and a nearly ideal scaling for **Ghost**. See also Figure 8.3 for a plot of these runtimes.

consider two different runs with process counts  $P_1$  and  $P_2$ , local element counts  $E_1$  and  $E_2$  and runtimes  $T_1$  and  $T_2$ . We compute the parallel efficiency  $e$  of the run with  $P_2$  processes in relation to the  $P_1$  run as the fraction

$$e_{\text{Ripple-balance}} = \frac{T_1 E_2}{T_2 E_1}. \quad (8.3)$$

As we see in Table 7.9 in the previous chapter, the runtime of **Ghost** depends linearly on the number of ghost elements per process. The number of ghosts is proportional to the surface area of a process's partition and thus ideally scales with  $\mathcal{O}((N/P)^{\frac{2}{3}})$ , with  $N$  the global number of elements [70]. Consider two runs with  $P_1$  and  $P_2$  processes as above and let  $G_1$  and  $G_2$  denote the numbers of ghost elements per process, then the parallel efficiency of the second run in relation to the first run is

$$e_{\text{Ghost}} = \frac{T_1 G_2}{T_2 G_1}. \quad (8.4)$$

We achieve ideal strong scaling efficiency for **Ghost** and even more than ideal efficiency for **Ripple-balance**, which hints at the runtime of **Ripple-balance** being slightly worse than  $\mathcal{O}(N)$ . We also observe that this basic variant of **Balance** is indeed by far the slowest of the AMR algorithms as we already hinted above.

### 8.3.3. Weak scaling

For weak scaling we increase the global number of elements while also increasing the process count, keeping the local number of elements nearly constant. Since with




Figure 8.3.: Strong scaling with tetrahedral elements. We show the runtimes of **Ripple-balance**, **Ghost**, and **Partition** for the test case from Section 8.3.1 with  $\ell = 8$ ,  $k = 2$  at time step  $t = 4\Delta t$ . The forest mesh consists of approximately  $1.9e9$  tetrahedra. We use 8,192 up to 131,072 processes on JUQUEEN with 16 processes per compute node. Ideally, **Ripple-balance** scales with the number of elements per process,  $N/P$ , and **Ghost** with the number of ghost elements per process,  $G/P$ . We show these measures scaled by a constant in black lines. As we observe in the plot and in Table 8.1, we achieve perfect scaling for **Ripple-balance** and nearly perfect scaling for **Ghost**. The runtime of **Partition** is below 0.1 seconds even for the largest process count.

each refinement level  $\ell$  the number of global elements grows by a factor of 8, we multiply the process count with 8 as well. We test the following configurations:

- Tetrahedral elements with 8,192 processes, 65,536 processes, and 458,752 processes, with refinement levels  $\ell = 8$ ,  $\ell = 9$ ,  $\ell = 10$ . This amounts to about 235k elements per process. Thus, the largest run has about  $107.8e9$  elements.
- Tetrahedral elements with 2,048 processes, 16,384 processes, and 131,072 processes, with refinement levels  $\ell = 8$ ,  $\ell = 9$ ,  $\ell = 10$ . Here we have about 155k elements per process, summing up to  $20.3e9$  elements on 131,072 processes.
- Hexahedral elements with the same process counts and level increased by one ( $162e9$  elements in total).

Note that 458,752 is actually 7 times 65,536. We choose it since it is the maximum possible process count on JUQUEEN with 16 processes per node, using all 28,672 compute nodes. The number of elements per process is thus about 14% greater than on the other process counts in the configuration. However, (8.3) and (8.4) still apply for computing the parallel efficiency.

The largest test case that we run is for hexahedra on 458,752 process with  $\ell = 11$  amounting to  $162e9$  elements.

We show the results in Table 8.2 and Figure 8.4. We notice that **Ghost** for tetrahedra is faster than **Ghost** for Hexahedra, which we explain by the smaller number

Tetrahedral case with $k = 2$ , $C = 0.8$ at $t = 4\Delta t$							
$P$	$\ell$	$E/P$	$G/P$	Ripple-balance		Ghost	
				Time [s]	Par. Eff.	Time [s]	Par. Eff.
8,192	8	234,178	17,946	687.0	100.0%	3.25	100.0%
65,536	9	233,512	18,282	732.5	93.5%	3.76	88.2%
458,752	10	266,494	20,252	913.5	85.5%	3.79	96.8%
2,048	7	117,630	10,999	305.5	100.0%	1.99	100.0%
16,384	8	117,089	11,311	336.2	90.4%	2.12	96.5%
131,072	9	116,756	11,478	360.8	84.0%	2.18	95.2%

Hexahedral case with $k = 2$ , $C = 0.8$ at $t = 2\Delta t$							
$P$	$\ell$	$E/P$	$G/P$	Ripple-balance		Ghost	
				Time [s]	Par. Eff.	Time [s]	Par. Eff.
8,192	9	309,877	34,600	268.5	100.0%	6.79	100.0%
65,536	10	310,163	35,136	272.8	98.5%	6.85	100.7%
458,752	11	354,746	38,833	313.0	98.2%	7.86	96.9%
2,048	8	156,178	21,536	131.4	100.0%	4.18	100.0%
16,384	9	155,702	22,036	132.7	98.7%	4.25	100.6%
131,072	10	155,460	22,284	134.4	97.2%	4.36	98.9%

Table 8.2.: Weak scaling for **Ripple-balance** and **Ghost** with tetrahedral (top) and hexahedral (bottom) elements. We increase the level by one and multiply the process count by eight to maintain the same number of local elements per process. Notice that the highest process count of 458,752 is only seven times 65,536 resulting in  $\sim 14\%$  more local elements. For hexahedra **Ripple-balance** has an overall better performance while **Ghost** performs better on tetrahedra, which we expect due to the lower number of faces per element. Both algorithms show good scaling with efficiencies greater than 85% (tetrahedra) and 96.9% (hexahedra). See also Figure 8.4. The maximum global number of elements is  $107.8e9$  with tetrahedra and  $162e9$  with hexahedra.

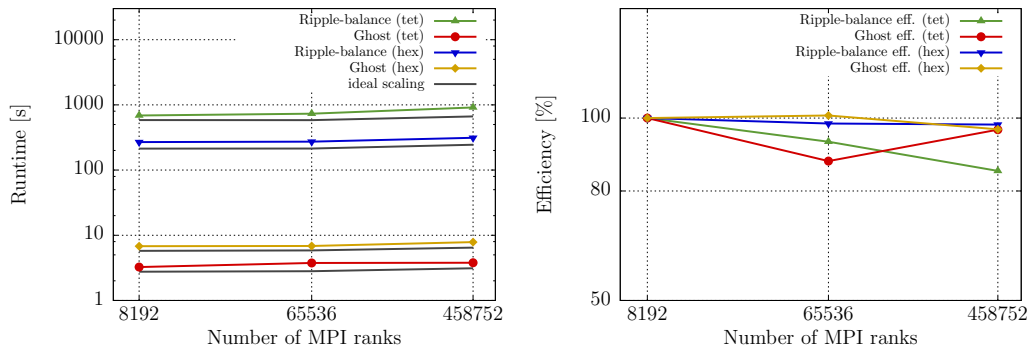


Figure 8.4.: Weak scaling results for tetrahedra with refinement levels 8, 9, and 10, and for hexahedra with refinement levels 9, 10, and 11. This amounts to about 233k elements per process for tetrahedra and 310k elements per process for hexahedra. This number increases for the 458,752 process runs, since 458,752 is only seven times 65,536, while we increase the number of mesh elements by the factor 8. On the left-hand side we plot the runtimes of **Ghost** and **Ripple-balance** with the ideal scaling in black. For **Ripple-balance** the ideal scaling line is based on the number of elements per process while for **Ghost** it is based on the number of ghost elements per process. On the right-hand side we plot the parallel efficiency in %. We display all values in Table 8.2

of ghosts due to less faces per element. **Ripple-balance** on the other hand is slower for tetrahedra than for hexahedra. The main part of **Ripple-balance** is a loop over all local elements. Thus, the overall slower performance of the tetrahedral Morton code compared to the cubical Morton code is probably the reason for this difference in runtime. We observe this behavior also in the weak parallel efficiency of **Ripple-balance** for tetrahedra which drops to 85.5% compared to the 98.2% efficiency for hexahedra. We even observe a strong scaling efficiency of more than 100% for tetrahedral **Ripple-balance**, which means that for tetrahedra the algorithm's runtime is not perfectly linear in the number of elements per process. This is also a hint that there is further potential to optimize **Ripple-balance**, which we expected, since there exist more efficient algorithms for 2:1 balancing on hexahedra [69].

For hexahedral **Ripple-balance**, and for hexahedral and tetrahedral **Ghost**, however, we observe excellent strong and weak scaling with efficiencies in the order of 95%.

## 9. A Numerical Application

In this chapter we discuss how an application could use the AMR routines described in this thesis. We implement a finite volume (FV) solver for the advection equation and discuss important data handling algorithms, such as communicating ghost data and interpolating data after changing the mesh.

We show that the solver works with different element types, in particular triangles and tetrahedra with the TM index described in Chapter 4 and quadrilaterals and hexahedra with the Morton index using the `p4est` implementation [25]. We also show runs on hybrid meshes consisting of triangles and quadrilaterals in 2D and of hexahedra, tetrahedra and prisms in 3D. For the low-level implementation of the prism elements we use the work from [81], which models a prism as the cross product of a line with Morton index and a triangle with TM-index.

### 9.1. The advection equation

We consider the  $d$ -dimensional advection equation. It describes how a quantity  $\phi$  is advected with a given flow  $u$  over time. For a compact domain  $\Omega \subset \mathbb{R}^d$  and a time-dependent flow function  $u: \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^d$ , we are interested in the solution  $\phi: \Omega \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$  of the PDE

$$\frac{\partial \phi}{\partial t} + \nabla \cdot (\phi u) = 0 \quad (9.1)$$

with initial condition

$$\phi(\cdot, 0) = \phi_0 \quad (9.2)$$

and appropriate boundary conditions. We assume that the flow  $u$  is divergence free, simplifying equation (9.1) to

$$\frac{\partial \phi}{\partial t} + u \cdot \nabla \phi = 0. \quad (9.3)$$

#### 9.1.1. Level-set functions

In this section we discuss level-set functions as a possible interpretation of  $\phi$ . There are various examples in which the advection equation is used to keep track of the movement—under the advection of the flow  $u$ —of a  $(d - 1)$ -dimensional interface between two disjoint subsets  $\Omega^1$  and  $\Omega^2$  with  $\Omega^1 \cup \Omega^2 = \Omega$  [3, 80, 99, 110, 111]. A typical example is two-phase flow [80, 138], where  $\Omega^1$  marks the region occupied by

the first phase of a fluid (i.e. liquid) and  $\Omega^2$  the region occupied by the second (i.e. gaseous). We pick the initial condition  $\phi_0$  such that it is continuous and satisfies

$$\phi_0(x) \geq 0, \text{ for all } x \in \Omega^1, \quad (9.4a)$$

$$\phi_0(x) < 0, \text{ for all } x \in \Omega^2. \quad (9.4b)$$

A common choice for  $\phi_0$  is the signed distance to the interface  $\Gamma = \overline{\Omega^1} \cap \overline{\Omega^2}$ , thus

$$\phi_0(x) = \begin{cases} \text{dist}(x, \Gamma), & x \in \Omega^1, \\ -\text{dist}(x, \Gamma), & x \in \Omega^2. \end{cases} \quad (9.5)$$

Let us define for each time step  $t$  the sets  $\Omega_t^i$  as the collection of all points in  $\Omega^i$  after they have been transported with  $u$ ,

$$\Omega_t^1 := \{x \in \Omega \mid \phi(x, t) \geq 0\}, \quad (9.6a)$$

$$\Omega_t^2 := \{x \in \Omega \mid \phi(x, t) < 0\}. \quad (9.6b)$$

Furthermore, we define  $\Gamma_t$  as the interface between  $\Omega_t^1$  and  $\Omega_t^2$ :

$$\Gamma_t := \overline{\Omega_t^1} \cap \overline{\Omega_t^2} = \{x \in \Omega \mid \phi(x, t) = 0\} = \phi^{-1}(0, t). \quad (9.7)$$

For a two-phase flow we interpret  $\Omega_t^i$  as the region of  $\Omega$  that is occupied by fluid  $i$  at time  $t$ . We can use the sign of  $\phi(x, t)$  to decide whether  $x \in \Omega_t^1$  or  $x \in \Omega_t^2$ .

**Remark 9.1.** Over the simulation time the approximated solution  $\phi$  may lose its signed distance property. A common observation is that the norm of the gradient at the zero level-set approaches zero, resulting in numerical irregularities when reconstructing the interface. A way to restore the signed distance property at a time step  $t$  is to replace  $\phi$  with a new level-set function  $\phi^*$  which is obtained as the solution to the Hamilton-Jacobi equation

$$\frac{\partial \phi^*}{\partial \tau} = \text{sign}(\phi_t)(1 - |\nabla \phi^*|) \quad (9.8)$$

with pseudo-time  $\tau$ . We omit this reinitialization process in our solver, since it does not affect the AMR routines. We refer the reader to [80, 98] for more details.

## 9.2. Numerically solving the advection equation

The numerical method we choose to solve the advection equation is a finite volume method with polynomial degree 0. This method leads to a rather simple solver that is easy to implement and provides first order convergence rates. We are well aware that far more accurate solvers exist. However, since we use the application as a proof-of-concept to show the coupling of application and AMR routines, we use degree 0 for the sake of simplicity. In fact, a cheap numerical solver is most



challenging for the AMR routines and their absolute runtimes. For a detailed description of the method, see [135] and the references therein. Although in practical applications the flow  $u$  could be the (discrete) solution of a fluid solver, we assume for our computations that  $u$  is given analytically.

Integrating equation (9.3) over a volume  $V \subset \Omega$  and applying the Gauss divergence theorem we obtain

$$\frac{\partial}{\partial t} \int_V \phi(x, t) dx + \int_{\partial V} \phi(s, t) u(s, t) \cdot \vec{n}(s) ds = 0, \quad (9.9)$$

with  $\vec{n}(s)$  being the outward pointing normal vector to  $V$  at  $s \in \partial V$ .

With the finite volume ansatz, we discretize the domain with mesh elements and consider equation (9.9) on each mesh element  $E$ . We model the approximate solution of  $\phi$  at time  $t$  with a constant value  $\phi_{E,t}$  on each element  $E$ . We interpret this value as an approximation to  $\phi$  at the midpoint  $m_E$  of  $E$ , thus  $\phi_{E,t} \approx \phi(m_E, t)$ . The initial condition at  $t = 0$  is then  $\phi_{E,0} = \phi_0(m_E)$ .

The boundary integral is an integral over the faces of  $E$ .

**Definition 9.2.** Let

$$\omega(E) := \{ (E', F) \mid E' \text{ is face-neighbor of } E \text{ across face } F \} \quad (9.10)$$

be the set of all pairs  $(E', F)$  of face-neighbors  $E'$  of  $E$  together with the respective face  $F$  of  $E$ .

Using this notation we obtain from equation (9.9) with  $V = E$ :

$$0 = \frac{\partial}{\partial t} \int_E \phi dx + \int_{\partial E} \phi u \cdot \vec{n} ds = \frac{\partial}{\partial t} \int_E \phi dx + \sum_{(E', F) \in \omega(E)} \int_F \phi u \cdot \vec{n} ds. \quad (9.11)$$

We now discretize the equation in time using a constant time step  $\Delta t$  and additionally discretize the right-hand side integrals as fluxes  $\psi(E, E'; F)$ . For the latter approximation we use an upwind method [44, 55]. To this end, let  $A(F)$  be the area of the face  $F$ . Furthermore, let  $\hat{u}$  be the flow  $u$  evaluated at the midpoint of  $F$ . We then define

$$\psi(E, E'; F) := \begin{cases} \phi_{E,t} (\vec{n} \cdot \hat{u}) A(F) & \text{if } \vec{n} \cdot \hat{u} \geq 0, \\ \phi_{E',t} (\vec{n} \cdot \hat{u}) A(F) & \text{if } \vec{n} \cdot \hat{u} < 0; \end{cases} \quad (9.12)$$

see also Figure 9.1, left. Note, that for a non-hanging face-connection we have

$$\psi(E, E', F) = -\psi(E', E, F). \quad (9.13)$$

We use the fluxes  $\psi$  and the approximation of  $\phi$  in equation (9.11) and discretize the time derivative in order to obtain

$$0 = \text{vol}(E) \frac{\phi_{E,t+\Delta t} - \phi_{E,t}}{\Delta t} + \sum_{(E', F) \in \omega(E)} \psi(E, E'; F) \quad (9.14a)$$

$$\Rightarrow \phi_{E,t+\Delta t} = \phi_{E,t} - \frac{\Delta t}{\text{vol}(E)} \sum_{(E', F) \in \omega(E)} \psi(E, E'; F). \quad (9.14b)$$

Thus, in each time step  $t + \Delta t$  we iterate over the (process local) elements and for each element  $E$  we iterate over all of its faces  $F$  with face-neighbors  $E'$  computing  $\psi(E, E'; F)$ . We then compute the value  $\phi_{E, t+\Delta t}$  according to (9.14) and continue with the next element. We store the values  $\psi(E, E'; F)$  in order to reuse them for the computation of  $\phi_{E', t+\Delta t}$ .

**Remark 9.3.** This iterating over the elements is a primitive version of the `Iterate` functionality. It remains a future project to implement a recursive version of `Iterate` as in [70] into `t8code`.

### 9.2.1. Hanging faces and face-neighbors

Since we use adaptive forests with non-conforming refinement methods, it is possible that an element  $E$  has more than one face-neighbor across a face  $F$ . Let  $E'$  be such a face-neighbor and let  $F'$  be the corresponding face of  $E'$ . Hence,  $F'$  is a subface of  $F$ . In this case, we must compute the face integral of (9.11) only over the part of the face of  $E$  that coincides with  $F'$ . Thus we set

$$\psi(E, E'; F) = -\psi(E', E; F') = \begin{cases} -\phi_{E', t}(\vec{n}' \cdot \hat{u}')A(F') & \text{if } \vec{n}' \cdot \hat{u}' \geq 0, \\ -\phi_{E, t}(\vec{n}' \cdot \hat{u}')A(F') & \text{if } \vec{n}' \cdot \hat{u}' < 0. \end{cases} \quad (9.15)$$

See the right-hand side of Figure 9.1 for an illustration.

Hence, given a leaf element  $E$  in the forest and a face  $F$  of it, we need to compute all face-neighbor leaf elements  $E'$  of  $E$  across the face  $F$ . This neighbor finding poses a challenge for adaptive forests when the face-neighbors of  $E$  are not of the same level  $\ell$  as  $E$ . If we assume that the forest is 2:1 balanced the possible refinement levels for face-neighbors restrict to  $\ell - 1$ ,  $\ell$ , and  $\ell + 1$ . In particular, all face-neighbors across the same face have the same refinement level. Using this, we can find the face-neighbors in the forest with the `t8_forest_half_face_neighbors` function from Section 7.2 and appropriate binary searches in the arrays of leaf elements and ghost elements.

**Remark 9.4.** Since we assume 2:1 balance of the forest to compute the face-neighbor leaves we call `Ripple-balance` after we adapt the forest to ensure the balance property.

### 9.2.2. The CFL number

In order for our explicit FV scheme to be numerically stable, we have to choose the time step  $\Delta t$  appropriately. If we choose it too big, the method becomes unstable. However, decreasing the time step increases the computational load, and thus increases the overall runtime. A well-known way to control the time step is the CFL number [43].




Figure 9.1.: Illustration of our notations for computing the fluxes as in (9.12) and (9.15). Left: A non-hanging face connection of a triangle element  $E$  across face  $F$ . Depending on the vector product  $\hat{u} \cdot \vec{n}$  we either use the value of  $\phi$  on  $E$  or on  $E'$  for the numerical flux and multiply with the area  $A(F)$  of the face  $F$ . Right: A hanging face connection of a triangle element  $E$  across face  $F$  (exploded view). In order to compute the flux across  $F$ , we sum up the two fluxes across the neighbor elements' faces  $F'$  and  $F''$ .

**Definition 9.5.** The **CFL-number** of an element  $E$  at time  $t$  is

$$C_{E,t} := u(m_E, t) \frac{\Delta t}{\text{vol}(E)^{\frac{1}{d}}}, \quad (9.16)$$

and the **global CFL-number** is

$$C_t = \max_E C_{E,t}. \quad (9.17)$$

In our solver, we choose an initial global CFL-number  $C$  and then compute  $\Delta t$  such that  $C_0 = C$ . We use this time step throughout the simulation, but it would be possible to change the time step during the simulation [72].

The FV theory states that the method is numerically stable for small values of  $C$ , i.e.  $C < 1$  [148]. In practice, we observe that we may need lower values in order to ensure stability. One reason is that  $\text{vol}(E)$  changes over time when the mesh is readapted.

### 9.2.3. The refinement criterion

In order to properly decide for a point  $x$  whether  $x \in \Omega_t^i$  it is important to resolve  $\Gamma_t$  as accurately as possible. This motivates refining the forest close to  $\Gamma_t$ , or rather its approximation  $\{E \mid |\phi_{E,t}| < \epsilon\}$ , as in [3]. Consider a leaf element  $E$  and let  $h = \text{vol}(E)^{\frac{1}{d}}$  be an approximation of its diameter and  $b > 0$  a given parameter, then we refine  $E$  at time step  $t$  if

$$|\phi_{E,t}| < bh. \quad (9.18)$$

The parameter  $b$  controls the width of the refinement band around  $\Gamma_t$ . We coarsen a family of elements if the equation does not hold for each of its members. Furthermore, we do not refine more than a specified maximum refinement level and do not coarsen above a minimum refinement level. This does not lead to element volume imbalances at the finest or coarsest level, since all trees have approximately the same volume.

#### 9.2.4. Error measurement

In our examples we use level-set functions with closed zero level-sets for the initial function  $\phi_0$ , for example a circle, or a sphere. Since the flow  $u$  is divergence free, the volume of the sets  $\Omega_t^i$  remains constant over time, if we do not have inflow or outflow. It is a typical behavior of FV solvers that we numerically loose or gain volume in  $\Omega_t^i$  over the computation. Whereby volume loss occurs in the convex region, which is the case for  $\Omega_t^2$  with our choices of initial conditions. Hence, we use the amount of lost volume in  $\Omega_t^2$  as a measurement for the computational error,

$$\mathcal{E}_t^{\text{vol}} := 1 - \frac{\text{vol}(\Omega_t^2)}{\text{vol}(\Omega_0^2)}. \quad (9.19)$$

We compute  $\text{vol}(\Omega_t^2)$  by summing up the volumes of those elements for which  $\phi_{E,t} < 0$ .

**Remark 9.6.** It is an inherent property of the method that, provided we have no in- or outflow, the integral of the level-set function

$$\int_{\Omega} \phi(t, x) dx \quad (9.20)$$

does not change over time (up to computational errors). This is due to the fact that the numerical flux (9.12) from element  $E$  to  $E'$  is exactly the negative as from element  $E'$  to  $E$ .

### 9.3. Handling application data

For each leaf element  $E$  we store some application specific data. In particular, this includes the value  $\phi_{E,t}$  of the approximated solution inside the element at the current time step. We store these data in an array which has one entry for each local leaf element and one for each ghost element. The data for the local elements is stored in order of their SFC index, while the order of the ghost elements is implicitly given by the `Ghost` algorithm.

In addition to the main mesh handling algorithms `New`, `Adapt`, `Partition`, `Ghost`, and `Balance` from the list in Section 2.5.2, we need subroutines in order to manage the application data. These perform the interpolation of data to a new forest after the mesh changes, the redistribution of the data after `Partition`, and the exchange of data of the ghost elements. These routines are common in AMR, see for example [28, 141], and we briefly summarize them in this section.

### 9.3.1. Interpolation

After we have computed the values  $\phi_{E,t}$  for each local element, we may modify the forest with **Adapt** to refine and coarsen the forest and **Balance** to reestablish the 2:1 balance. Thus, the current forest  $\mathcal{F}$  is changed to a new forest  $\mathcal{F}'$ . For all new elements in  $\mathcal{F}'$  we need to calculate an interpolated value of  $\phi$ , which we do in the following way as described in [28]. We restrict to non-recursive refinement. Hence, all elements in  $\mathcal{F}'$  result from an element in  $\mathcal{F}$  by either refining once, coarsening once, or keeping the element as it is. In the first case, an element  $E$  in  $\mathcal{F}$  is refined into  $n$  children  $E_i$ ,  $0 \leq i < n$ , and we set the value of the new elements to the one of the parent:

$$\phi_{E_i,t} = \phi_{E,t} \quad (9.21)$$

for all  $0 \leq i < n$ . In the second case, a family  $\{E_i\}$  of  $n$  elements is coarsened into their parent  $E$ . We assign the average of the values of the  $E_i$ , thus

$$\phi_{E,t} = \frac{1}{n} \sum_i \phi_{E_i,t}. \quad (9.22)$$

If an element is unchanged, we also do not change the value of  $\phi$ .

The challenge is to identify the corresponding pairs of new and old elements. In order to do so, we iterate simultaneously through the leaf elements of  $\mathcal{F}$  and  $\mathcal{F}'$  with indices  $i$  and  $j$ , starting with  $i = j = 0$ . Let  $l_i$  and  $l_j$  be the refinement levels of the  $i$ -th element in  $\mathcal{F}$  and the  $j$ -th element in  $\mathcal{F}'$ . If  $l_i = l_j$ , they are the same element, hence we do not change the  $\phi$  value and increase both indices by 1. If  $l_i = l_j - 1$ , we know that the element was refined into  $n$  children. We carry out the interpolation (9.21) and increment  $i$  by 1 and  $j$  by  $n$ . If  $l_i = l_j + 1$ , the element is the first in a family of  $n$  elements that was coarsened and we compute the average (9.22). We then increment  $i$  by  $n$  and  $j$  by 1.

We refer to this routine as **Interpolate**.

### 9.3.2. Repartitioning of data

After the forest is adapted and the new element values are interpolated, we may repartition the forest via **Partition** in order to maintain a balanced load. After **Partition** the element data needs to be partitioned as well. Partitioning the element data follows the same logic as partitioning the elements. We call the routine **Partition\_data**.

### 9.3.3. Ghost exchange

Before every new time step, we need to update the function values in the ghost elements. Thus, each process has to send the values  $\phi_{E,t}$  for its inter-process boundary leaf elements to all processes that have face-neighbors of this boundary element. This operation is called **Ghost\_exchange** [30].

Its input is the data array with valid entries  $\phi_{E,t}$  for all process local leaf elements. On output the values in the data array corresponding to the ghost elements are filled with the entry of the respective owner process.

The implementation is straight-forward if we store the indices of boundary elements and of which remote processes these are ghosts at the time of creation of the ghost layer, hence during `Ghost`.

## 9.4. Tests on a unit cube geometry

We run tests in 2D and 3D with  $\Omega = [0, 1]^d$ , the  $d$ -dimensional unit cube with periodic boundary conditions for  $\phi$ .

### 9.4.1. The 2D test case

For the 2D tests, we use a flow  $u$  that simulates a rotation around the midpoint  $(0.5, 0.5)$  of  $\Omega$ ,

$$u(x, y) = 2\pi \begin{pmatrix} y - 0.5 \\ -(x - 0.5) \end{pmatrix}. \quad (9.23)$$

This  $u$  is divergence free and chosen such that for the analytical solution at time  $t = 1$  we obtain

$$\phi(\cdot, 1) = \phi_0. \quad (9.24)$$

As level-set function  $\phi_0$  we choose the signed distance to a circle of radius 0.25 with midpoint  $(0.6, 0.6)$ , hence

$$\phi_0(x, y) = \sqrt{(x - 0.6)^2 + (y - 0.6)^2} - 0.25. \quad (9.25)$$

We compare three different coarse meshes of  $[0, 1]^2$ , which we show in the top part of Figure 9.2. The first one consists of one quadrilateral tree; for the second one we divide the square along a diagonal into two triangle trees; for the third one we use a hybrid coarse mesh of four triangle trees and two quadrilateral trees. For the refinement levels our solver accepts two arguments:  $\ell$  and  $r$ .  $\ell$  describes the minimum refinement level and  $r$  the maximum number of additional refinement levels. Thus, we start with a uniform level  $\ell$  forest and refine it up to level  $\ell + r$  before starting the computation. We depict adaptive forests with an initial adaptive refinement along the zero-level set of  $\phi_0$  in Figure 9.2.

We show a computation of the hybrid forest with  $\ell = 4$  and  $r = 3$  in Figure 9.3. This figure also shows the flow field  $u$ .

### 9.4.2. The 3D test case

For computations on the 3D unit cube geometry, we construct a specific flow function.

**Definition 9.7.** Let  $f \in C^2([0, 1])$  with  $f(0) = f(1) = 0$ . We define the vector field  $u^f : [0, 1]^3 \rightarrow \mathbb{R}^3$  as

$$u_1^f(x, y, z) = f(x)(f'(y) - f'(z)) \quad (9.26a)$$

$$u_2^f(x, y, z) = -f'(x)f(y) \quad (9.26b)$$

$$u_3^f(x, y, z) = f'(x)f(z). \quad (9.26c)$$

$u^f$  has the following properties:

**Lemma 9.8.**  $u^f$  is divergence free and has no outflow from the unit cube, thus  $u^f \cdot \vec{n} = 0$  at each face of  $[0, 1]^3$ .

*Proof.* The first part follows from computing the partial derivative of  $u^f$  and the second part follows since

$$u_1^f(0, y, z) = u_2^f(x, 0, z) = u_3^f(x, y, 0) = 0, \quad \text{and} \quad (9.27a)$$

$$u_1^f(1, y, z) = u_2^f(x, 1, z) = u_3^f(x, y, 1) = 0. \quad (9.27b)$$

□

For our test case we use  $f(x) = \sin(\pi x)$  and consider the time-dependent flow

$$u : [0, 1]^3 \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^3, \quad (9.28a)$$

$$(x, y, z, t) \mapsto \begin{cases} u^f(x, y, z), & t < 0.5, \\ -u^f(x, y, z), & t \geq 0.5. \end{cases} \quad (9.28b)$$

In analogy to the 2D case, the analytical solution to the advection equation satisfies  $\phi(\cdot, 1) = \phi_0$ , which allows us to compute the numerical error at time  $t = 1$ . We show the flow  $u^f$  in Figure 9.4.

For the initial level-set function  $\phi_0$  we use a signed distance to a sphere of radius 0.25 and midpoint  $(0.6, 0.6, 0.6)$ , hence

$$\phi_0(x, y, z) = \sqrt{(x - 0.6)^2 + (y - 0.6)^2 + (z - 0.6)^2} - 0.25. \quad (9.29)$$

We use three different coarse meshes of  $[0, 1]^3$ . Firstly, one singular hexahedral tree; secondly, six tetrahedral trees as in Figure 4.2; and thirdly, a hybrid mesh of six tetrahedral trees, six prism trees, and four hexahedral trees, which we depict in Figure 9.4. Figure 9.5 additionally shows a detailed cut out view of the hybrid mesh, displaying refined forest elements within the unit cube.

### 9.4.3. Convergence tests

We perform uniform and adaptive runs for quadrilaterals, triangles, hexahedra, tetrahedra, and hybrid meshes on  $\Omega = [0, 1]^d$  ( $d$  being the dimension) with small values of  $\ell$  and  $r$  in order to verify the convergence of our method in terms of  $\mathcal{E}_t^{\text{vol}}$ .




Figure 9.2.: Adapting three different unit square meshes at a circle with midpoint  $(0.6, 0.6)$  and radius  $0.25$ , modeled as the zero level-set of a function  $\phi$ . We show the coarse mesh in the top row and the adapted forest mesh in the bottom row at time  $t = 0$ . We model the unit square using one quadrilateral tree (left, red), two triangle trees (middle, blue), and as a hybrid mesh of four triangular and two quadrilateral trees (right). Starting from a uniform forest, we use six refinement levels. The initial uniform forest has level 2 in the quadrilateral and triangle case, and level 1 in the hybrid case. In the bottom row the color represents the values  $\phi_{E,t}$  of the approximated level-set function.






Figure 9.3.: A simulation of the 2D test case on the 2D hybrid mesh. Top left: The rotational flow  $u$  from (9.23). Top right, bottom left, bottom right: Times  $t = 0$ ,  $t = 0.5$ , and  $t = 1$  of the simulation with initial refinement level  $\ell = 4$ , and with  $r = 3$  adaptive refinement levels. The maximum refinement level is 7 and there are about 20,000 elements. The number of elements in an equivalent uniform level 7 forest is 98,304. We depict the zero-level set of the analytical solution in white and the zero-level set of the computed solution in purple.




Figure 9.4.: Illustrations of our 3D test case. Top row: The 3D hybrid coarse mesh of the unit cube consisting of six tetrahedra (blue), six prisms (orange), and four hexahedra (red). The right-hand side picture shows an exploded view of the coarse mesh. Bottom left: Streamlines of the flow  $u^f$  from (9.28) with  $f(x) = \sin(\pi x)$ ; the colors serve the purpose of distinguishing the streamlines. Bottom right: Some flow vectors  $u^f(x)$ ; color and size of the arrows indicate the magnitude  $\|u^f(x)\|$  of the flow. We observe that the flow rotates around a diagonal line through the cube and that there is no outflow, as shown by Lemma 9.8.




Figure 9.5.: A view inside the hybrid 3D mesh modelling the unit cube (cf. Figure 9.4). The picture shows the region near the center of the unit cube where tetrahedra (blue), prisms (orange), and hexahedra (red) are close to each other. We see in detail how the prisms serve to transition from tetrahedra to hexahedra. We use parameters  $\ell = 2$  and  $r = 5$  to generate this picture, thus the finest elements are of refinement level 7.

For each coarse mesh type we fix a CFL number  $C$ . Hence, each time we increase the level by one,  $\Delta t$  is divided by two, doubling the number of time steps. We expect a first order convergence of the error at the final time. In Table 9.1 we display the results for the 2D and 3D test cases above and we indeed observe a first order convergence rate. We additionally display the number of elements, which is constant over all time steps in the uniform case and an average over all time steps for the adaptive runs. We see that with the adaptive runs we may achieve approximately the same error as with uniform runs of the same maximal level while using significantly less mesh elements.

#### 9.4.4. Large scale tests

In this section We investigate the runtimes and scaling behavior of the solver with the 3D tests with tetrahedra on the JUQUEEN supercomputer [73]. For all of the following runs we use 16 MPI ranks (processes) per compute node, thus 1 per compute core.

Firstly, we compare the amount of runtime spent in the AMR routines with the actual time spent using the FV solver. To this end, we perform a strong scaling tests with 16,384 processes, 32,768 processes, and 65,536 processes and an average number of tetrahedra over all time steps of 8,344,140. We display the results in Figure 9.6, where we show the runtimes of the different AMR routines and their respective proportion of the overall runtime. We exclude the runtime of the `Ripple-Balance` routine here, since it is non-optimized and would distort the results. We notice that the percentage of time spent in the AMR routines increases with decreasing number of elements per process from around 35% to 50%, which is expected for strong scaling with small numbers of elements per process. Using a more involved numerical solver—for example higher order FV or DG—would certainly decrease these percentages, since then the amount of solver time per mesh element would increase while the time spent in AMR routines would remain the same.

Secondly, we perform strong scaling tests with 131,072 processes, 262,144 processes, and 458,752 processes, each case consisting of about  $2.3e9$  mesh elements. The equivalent uniform mesh would have more than  $50e9$  elements. We use a short simulation time of  $T = 0.0005$  to keep the overall runtimes below 100 seconds. We display the scaling results in Figure 9.7—not counting the runtime of the non-optimized `Ripple-Balance`—split up into runtime of the AMR routines only and total runtime (AMR + solver). Furthermore, we list the exact runtimes in Table 9.2 and compute the parallel efficiency. Compared to the base-line run with 131,072 processes we obtain an ideal strong scaling efficiency on 458,752 processes if we do not take `Ripple-Balance` into account. Including `Ripple-Balance`, the efficiency is still above 90% for 262,144 processes, and 86.2% (AMR), respectively 89.6% (total), for 458,752 processes. We conclude that providing a state-of-the-art `Balance` has the potential to increase the scaling behavior to the full 100%.

2D quadrilaterals, $C = 0.5$				2D triangles, $C = 0.1$				2D hybrid, $C = 0.1$			
$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$	$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$	$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$
4	0	256	95.7%	4	0	512	70.4%	3	0	384	93.8%
5	0	1,024	60.2%	5	0	2,048	41.2%	4	0	1,536	60.9%
6	0	4,096	33.7%	6	0	8,192	22.5%	5	0	6,144	34.1%
7	0	16,384	18.2%	7	0	32,768	11.4%	6	0	24,576	18.7%
3	3	3,613	34.0%	3	3	5,358	23.5%	2	3	3,948	35.8%
4	3	7,475	21.9%	4	3	11,542	15.0%	3	3	8,323	24.9%

3D hexahedra, $C = 0.25$				3D tetrahedra, $C = 0.1$				3D hybrid, $C = 0.1$			
$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$	$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$	$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$
4	0	4,096	94.2%	4	0	24,576	70.3%	3	0	8,192	93.4%
5	0	32,768	63.2%	5	0	196,608	41.0%	4	0	65,536	63.7%
6	0	262,144	34.5%	6	0	1,572,864	21.8%	5	0	524,288	34.7%
7	0	2,097,152	18.1%	7	0	12,582,912	11.3%	6	0	4,194,304	18.1%
3	3	113,851	34.8%	3	3	652,615	22.0%	2	3	173,635	35.1%
4	3	393,079	19.44%	4	3	2,305,350	12.1%	3	3	559,860	19.8%

Table 9.1.: We verify convergence of the advection solver in terms of volume loss of  $\Omega^2$  at time  $t = 1$ . We test six different coarse meshes: 2D quadrilateral (top left), 2D triangles (top middle), 2D hybrid (top right), 3D hexahedra (bottom left), 3D tetrahedra (bottom middle), and 3D hybrid (bottom right); see also Figures 9.4 and 9.2. In each case the domain  $\Omega$  is the unit cube of the corresponding dimension. In 2D we use as flow  $u$  a rotation around the center of the cube from equation (9.23), see also the top left of Figure 9.3; in 3D we use the flow  $u$  from equation (9.28) with  $f(x) = \sin(\pi x)$ ; see Figure 9.4 for an illustration. In the top part of each table, we list a uniform level  $\ell$  test with increasing values of  $\ell$ . We fix the CFL number  $C$ , which results in doubling the number of time steps each time the level increases. In the bottom part of each table, we show the results for adaptive refinement to level  $\ell + 3$ . For the refinement criterion (9.18) we use a band width of  $b = 4$ . In each case we verify a first order convergence rate.

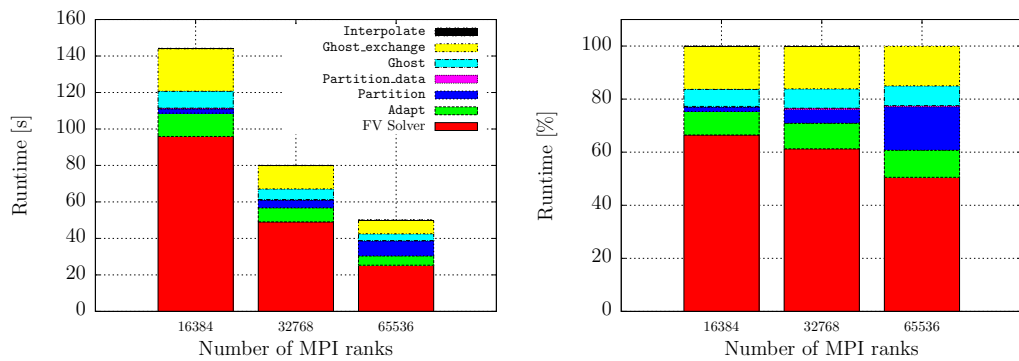


Figure 9.6.: Strong scaling with tetrahedral elements. Using  $\ell = 4$  and  $r = 4$ , the maximum refinement level is 8. The CFL number is  $C = 0.1$  and the band-width parameter for refinement is  $b = 4$ . The average count of elements is 8,344,140. Left: Total runtime of the different AMR routines and the solver. Right: Relative runtimes of the same methods. We observe that with decreasing number of elements per process the relative runtime of AMR routines increases from approximately 35% to 50%. Since the CFL number is 0.1 we only change the mesh in every 10-th time step. Thus, `Adapt`, `Partition`, `Partition_data`, `Ghost`, and `Interpolate` are only called in these time steps, while `Ghost_exchange` is called in every time step, explaining the relatively large portion of runtime taken up by `Ghost_exchange`.

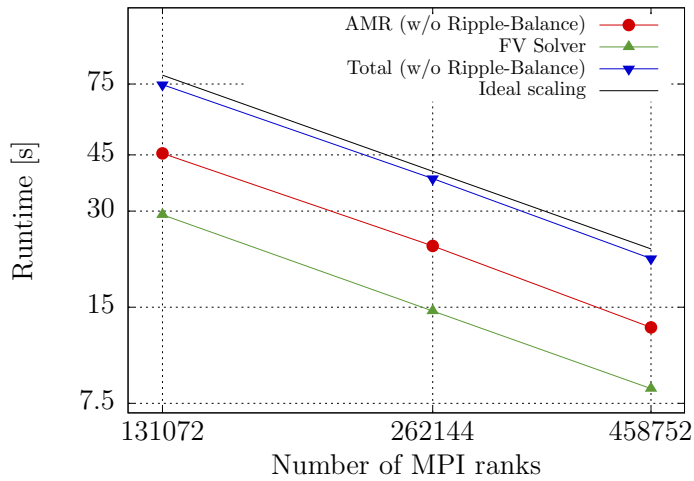


Figure 9.7.: Strong scaling results with tetrahedra with 131,072 processes, 262,144 processes, and 458,752 processes. We use  $\ell = 5$ ,  $r = 6$ , resulting in  $2.3e9$  mesh elements. In order to decrease the overall runtime, we pick a simulation end time of  $T = 0.0005$  and adapt every 10-th time step. We show the runtime split between the FV solver (green triangles pointing up) and the AMR routines (red dots), as well as the total runtime (blue triangles pointing down). The AMR routines are `Adapt`, `Partition`, `Ghost`, `Ghost_exchange`, `Partition_data`, and `Interpolate`. The black line represents the ideal scaling behavior. We reach this ideal behavior on 458,752 MPI ranks with a parallel efficiency of 100% compared to the baseline with 131,072 MPI ranks. See also Table 9.2 where we explicitly list the runtimes.

$P$	Runtime (without Ripple-Balance)		par. efficiency	
	AMR	Total	AMR	Total
131,072	45.5s	74.7s	–	–
262,144	23.3s	37.9s	97.6%	98.5%
458,752	13.0s	21.3s	100.0%	100.0%

$P$	Runtime (including Ripple-Balance)		par. efficiency	
	AMR	Total	AMR	Total
131,072	75.1s	104.4s	–	–
262,144	41.4s	56.0s	90.7%	93.2%
458,752	24.9s	33.3s	86.2%	89.6%

Table 9.2.: In this table we present the runtimes and compute the parallel efficiency of the AMR routines and the total runtime of the advection solver for the tetrahedral strong scaling test on 131,072 up to 458,752 processes. See also Figure 9.7. The top part of the table shows the runtimes without **Ripple-Balance**, we obtain a parallel efficiency of 100% for the run on 458,752 processes compared to the base line run with 131,072 processes. In the bottom part, we include **Ripple-Balance**. The overall parallel efficiency is then 86.2% for the AMR routines and 89.6% for the total runtime.

#### 9.4.5. Comparison to uniform meshes

It is sensible to ask whether we profit from using AMR at all. When we omit the adaptive refinement and use the initial level  $\ell$  uniform mesh in the solver, there is no overhead due to AMR algorithms and all of the compute time is used to solve the actual numerical equation. However, to reach the same accuracy, more mesh elements are needed, which in turn results in a higher memory usage and possibly larger overall runtime.

In Table 9.3 we compare adaptive and uniform runs on tetrahedral meshes with 32,768 and 65,536 processes. In the uniform case we use a refinement level of  $\ell = 8$  and in the adaptive tests we use an initial level  $\ell = 4$  mesh which we adapt  $r = 4$  further levels, such that the finest elements are of level 8 as well. First of all we notice that the adaptive meshes only use 8.3% as many elements as the uniform meshes (8.3e6 and 100.6e6 elements) while resulting only in a slightly larger computational error of 7.6% volume loss compared to the 5.8% volume loss in the uniform case. This large reduction in the number of elements points to a significant decrease in memory usage.

Furthermore, we observe that the adaptive runs need less than half the runtime as the uniform runs. These are total runtimes that include the relatively slow **Ripple-balance** routine and thus we can expect that the overall gain would be even better with an optimized **Balance** routine. For applications with more highly localized physics, adaptivity may reduce the number of elements in relation to uni-



$P$	$\ell$	$r$	mesh size	$\mathcal{E}_1^{\text{vol}}$	time steps	Runtime
32,768	8	0	100,663,296	5.8%	5,196	469.9s
32,768	4	4	8,336,500	7.6%	3,445	159.3s
65,536	8	0	100,663,296	5.8%	5,196	240.3s
65,536	4	4	8,339,130	7.6%	3,445	102.3s

Table 9.3.: We compare adaptive and uniform runs of the same problem, once with 32,768 and once with 65,536 processes. The finest refinement level in all cases is 8. The adaptive runs only use 8.3% as many mesh elements as the uniform runs and need less than half the runtime to obtain an only slightly larger computational error. The runtimes in the last column are the total runtimes of the solver and in particular include the non-optimized `Ripple-balance` routine in the adaptive cases.

form meshes by over three orders of magnitude, which will lead to proportionally higher savings.

#### 9.4.6. A test with a larger coarse mesh

We close this section with an application of the solver on a non-trivial domain with a medium sized coarse mesh. We use the example of two-dimensional potential flow around a disk [12]. This flow is an analytical solution to the flow of an incompressible fluid without viscosity around a disc with radius  $R$  and midpoint at the origin, with the flow being constant 1 in  $x$ -direction and 0 in  $y$ -direction far outside of the disc. In polar coordinates the flow field  $u$  is given by

$$u(r, \phi)_r = \left(1 - \frac{R^2}{r^2}\right) \cos(\phi), \quad (9.30a)$$

$$u(r, \phi)_\phi = \left(-1 - \frac{R^2}{r^2}\right) \sin(\phi). \quad (9.30b)$$

As radius we chose  $R = 0.15$ , and we illustrate the flow in Figure 9.8.

Since the flow is symmetric around the  $x$ -axis, we restrict our attention to the region with  $y$ -coordinates greater or equal zero. As domain  $\Omega$  we choose the rectangle  $[-0.5, 1] \times [0, 0.75]$  with the disk cut out.

We model this domain with a hybrid coarse mesh of 238 triangles and 351 quadrilaterals as in Figure 9.9. The purpose of the quadrilaterals is to properly resolve the flow close to the curved boundary. In order to do so, we use a boundary layer of thin quadrilaterals that are stretched in the direction of the flow, which we display in Figure 9.9 on the right.

**Remark 9.9.** We modify our refinement criterion. It is still sensible to refine in a band around the zero level-set as in (9.18). However, the sizes of the coarse mesh elements differ largely. In particular the quadrilateral elements at the circle boundary are very small compared to the triangular elements filling the rest of the




Figure 9.8.: Streamlines (white) of the flow (9.30) around a disk with radius  $R = 0.15$ . The arrows on the left-hand side indicate the inflow vector (orange) while the background color indicates the magnitude of the flow velocity.




Figure 9.9.: Left: We model the rectangle  $[-0.5, 1] \times [0, 0.75]$  with the disk of radius 0.1 cut out as a hybrid triangle/quadrilateral mesh with 238 triangles and 351 quadrilaterals. Right (zoomed in): We use the quadrilaterals (red) to resolve the flow close to the curved boundary and triangles (blue) to mesh the remaining domain. This mesh was created with `Gmsh` [57]. Throughout the mesh the sizes of the coarse elements differ by several orders of magnitude, which motivates taking the volume of an element into account when refining according to (9.18).

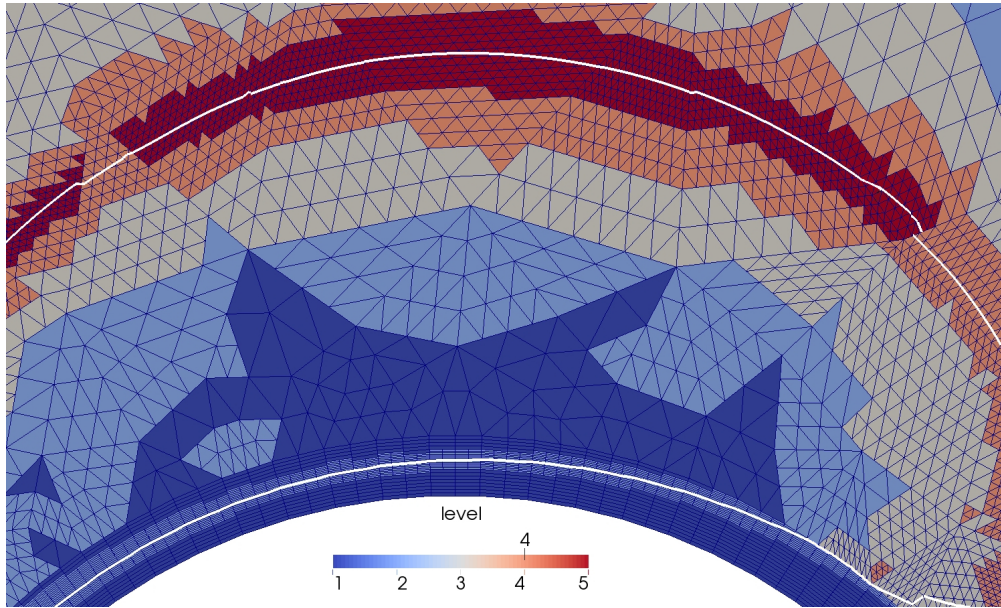


Figure 9.10.: Zoomed in: The different refinement levels close to the zero level-set (white). We refine an element if it is close to the zero level-set and its volume is larger than a given lower bound. Thus, the finest elements close to the zero level-set all have comparable volumes. The difference in refinement levels is observed when we compare the level 5 triangles (red) in the top part of the image with the level 2 quadrilaterals (light blue) in the bottom part.

domain. In the final refined mesh, all elements close to the zero level-set should have approximately the same size. In order to achieve this, we refine an element if it fulfills criterion (9.18) only if its volume is above a certain threshold, which is determined by the volume of the smallest elements in the mesh; see Figure 9.10 for an illustration. We thus replace the coarsest and finest level restriction described in Section 9.2.3 by a largest and smallest volume restriction.

As initial level-set function  $\phi_0$  we choose the signed distance function to a circle with radius 0.1 and midpoint  $(0.2, 0.11)^t$ . In particular, this means that the zero level-set is advected close to the curved boundary.

In Figure 9.11 we show six different time steps of the computation with initial uniform level  $\ell = 1$  and  $r = 4$  adaptive refinement levels. It is clearly visible how close the zero level-set is to the circular hole.




Figure 9.11.: Six different time steps of the numerical solution to the advection of a circular interface with the flow (9.30) and maximum refinement level 5. The interface moves very close to the circular domain boundary. This movement is well resolved since we use flat quadrilaterals close to this boundary. In the more orange region to the left of the zero level-set we observe the effect that  $\phi$  loses its signed distance property, which we addressed in Remark 9.1. This results in elements keeping a finer level than necessary.

## 10. Conclusion

In this thesis we are concerned with the development of a new space-filling curve (SFC) for parallel simplicial adaptive mesh refinement, as well as scalable data structures and algorithms for parallel tree-based adaptive mesh refinement (AMR) with general element types (such as for example triangles, tetrahedra, quadrilaterals, hexahedra, and prisms) and support for hybrid meshes.

Existing libraries for tree-based AMR focus on a single type of element and SFC, such as for example `p4est` [34] for quadrilaterals/hexahedra with the Morton SFC. A major contribution of this thesis is the development of an element-type independent approach by separating the high-level AMR algorithms from the low-level element-type implementation. By modifying and restructuring the existing meshing algorithms we extend the AMR functionality to arbitrary element types. In particular, to the best of our knowledge, we present the first discussion and applications of tree-based AMR on tetrahedral and hybrid meshes.

While we applied our concept with little effort to mesh refinement and coarsening, as well as partitioning, significant work was necessary to establish the construction of a layer of ghost elements. Additionally, we devised a new communication reducing algorithm for coarse mesh partitioning.

Besides a thorough theoretical investigation, we implemented the `t8code` AMR library to demonstrate practicability and verified strong and weak scaling to hundreds of thousands of parallel processes on current supercomputing systems. To confirm that our algorithms provide all features required for a basic numerical application, we implemented a finite volume solver for the advection equation. We find that our implementation supports face-based numerical applications in a modular and non-intrusive way.

A starting motivation for this thesis was to find an SFC for triangles and tetrahedra that has similar properties to the Morton SFC for quadrilaterals and hexahedra, such as fast computation via bitwise interleaving and constant-time, level-independent low-level algorithms (constructing children, face-neighbors, parents, etc.). Though it was not certain that such an SFC exists, we were able to find a suitable construction leading to the tetrahedral Morton (TM-) SFC. In order to transfer the bitwise interleaving concept to simplices, we introduced the type of a simplex as additional information, and used the types of an element's ancestors to interleave with the anchor node coordinates. Despite its similarities to the Morton SFC, segments of the TM-SFC can produce more connected components when partitioning the mesh in parallel, possibly increasing the amount of parallel communication. We investigated this further and proved upper bounds for the possible

counts and demonstrated that in practical applications less than 7% of the segments have more than three connected components. Throughout this thesis we observed that the performance of the TM-SFC is comparable to that of the Morton SFC. Our discussions and examples show that the TM-SFC is an excellent choice for an SFC for triangles and tetrahedra and is a competitive addition to the multitude of existing SFCs.

In order to model complex domain geometries with arbitrarily large numbers of trees, we investigated parallel partitioning of the coarse mesh. We devised a new approach that reduces parallel communication while maintaining the load-balance of the fine mesh. Because we duplicate a tree's data on all processes that have local elements of this tree, repartitioning requires the solution of an  $n$ -to- $m$  communication problem. We developed a sophisticated repartitioning algorithm to ensure that despite information being duplicated, no data is sent multiple times. Additionally, we incorporated the communication of ghost trees and demonstrated scaling on up to 917e3 processes with tetrahedral and hexahedral tree types and meshes of up to 371e9 trees with faster execution times than the fine mesh partitioning.

For our development of element-type independent high-level AMR algorithms, the `Ghost` algorithm to create a layer of ghost elements at each process turned out to be the most challenging. The most complex step in this algorithm is constructing the face-neighbors of a given element for which we formulate a new element-type independent algorithm. With our approach we gain maximum flexibility with regards to the implementations of the SFCs used in the neighboring trees. We optimized our `Ghost` routine by using recently developed recursive search routines for tree-based AMR. This algorithm has optimal strong and weak scaling behavior and we obtain runtimes comparable to current state-of-the-art implementations with fixed element-types [70].

We concluded our investigation of element-type independent AMR by showing that we can implement a scalable straightforward version of the 2:1 Balance routine based on a ripple algorithm [143, 144], using the available functionalities such as `Ghost`.

By implementing and investigating a finite volume solver we demonstrated that `t8code` can easily be integrated in numerical solvers and maintains excellent runtimes and scaling behavior in a practical application setting. We underlined the flexibility in terms of element-types by performing tests on different types of meshes including 2D and 3D hybrid meshes.

Concluding, we developed a new SFC for simplicial mesh refinement and a complete software library for element-type independent tree-based AMR that supports different types of elements as well as hybrid meshes. We demonstrated excellent strong and weak scaling behavior on current high-performance computing systems and showed that our software can be directly used by numerical applications.

Our modular approach of separating the high-level and low-level algorithms proved to be very successful and leads to a strong flexibility, allowing us to reuse the same high-level algorithms for different element types and SFCs. We demonstrate this for

example by testing `Ghost` and `Balance` with hexahedral and with tetrahedral elements and by applying the advection solver to various different meshes including 2D and 3D hybrid meshes. With the definition of an API for the low-level algorithms, we enable users of `t8code` to reuse existing implementations for an element-type, such as we did throughout this thesis by using the `p4est` implementation for the quadrilateral and hexahedral Morton SFC. This can greatly reduce the amount of work when extending existing applications with the tree-based AMR functionality from `t8code`. Another advantage is that additional element-types can be added to `t8code` with little effort, such as shown in a Bachelor's thesis that includes the support for prisms into `t8code` without any necessary changes to the high-level algorithms [81].

We are hopeful that the techniques presented in this thesis can lead to tree-based AMR being used in areas that are currently dominated by unstructured meshing, such as large-scale engineering applications with hybrid meshes; for example [77, 95, 147].



## 11. Outlook

Our results and algorithms provide a significant improvement in tree-based AMR techniques, especially considering their generality and flexibility. Nevertheless, there are various possibilities for further research and improvement, of which we briefly discuss some examples here.

A first possible extension is to develop an SFC for pyramidal refinement and establish the necessary low-level algorithms. This extends the functionality to fully hybrid meshes consisting of tetrahedra, hexahedra, prisms, and pyramids and thus allows for complex engineering applications as for example described in [77, 147]. A similar possible extension are 4D elements such as hypercubes and 4D simplices. These are for example used to model 4D space-time [54].

Also, if we consider the history of development of tree-based quadrilateral/hexahedral AMR, with particular emphasis on the `p4est` library, we see several future research opportunities:

As we discuss in Sections 8 and 9, `Balance` is the most expensive high-level algorithm in `t8code`, which was also the case for the original implementation in `p4est`. Thus, developing a state-of-the-art version along the lines of [69] would significantly reduce the runtime and improve the overall scaling behavior of `Balance`. A particular challenge in this development will be the generalization of an element's insulation layer to arbitrary element types and across trees of different types.

A second high-level algorithm to improve is `Iterate` to execute a user provided callback on all mesh elements and element-to-element interfaces. In this thesis we use a basic loop; however, a sophisticated implementation using search algorithms that handles hanging nodes automatically is possible. Such a universal mesh topology iterator for quadrilaterals/hexahedra is described in [70], and we believe these techniques can be transferred to the general approach presented in this thesis.

Other points are the extension of the geometry handling to edge and vertex neighbors and the development of a vertex numbering scheme [70]. This can then be utilized by application codes that are vertex based, such as for example finite element solvers. A starting point could be the implementation of the `Nodes` algorithm from [34], which would require handling of hanging edges and the computation of edge neighbors as a low-level function and in the coarse mesh.

Furthermore, it would be interesting to use higher order geometry representations for the coarse mesh. Each tree in the coarse mesh is equipped with geometry interpolation points, which allows to use curved geometries for each tree; see for example [150, 154]. This technique can improve the overall geometric accuracy while decreasing the number of coarse mesh trees. An implementation into `t8code` would be straightforward, utilizing the existing data management algorithms of the



coarse mesh.

## A. The Low-Level API

In this appendix we list all low-level functions from the `t8code` library. In order to introduce a new element type, all of these functions have to be implemented for this element type. Currently, `t8code` supports the following element types.

- 0D Vertices
- 1D Lines with Morton order
- 2D Quadrilaterals with Morton order using the `p4est` library [25]
- 3D Hexahedra with Morton order using `p4est`
- 2D Triangles with Tetrahedral-Morton order
- 3D Tetrahedra with Tetrahedral-Morton order
- 3D Prisms as a cross product of Lines and Triangles [81]

We list all low-level functions together with a brief description of their effect in Table A.1. For more details, we refer the reader to the documentation of `t8code` [139].

Function name	effect
<code>t8_element_maxlevel</code>	Return the maximum possible refinement level
<code>t8_element_child_eclass</code>	Return the element type of the $i$ -th child
<code>t8_element_level</code>	Return an element's level
<code>t8_element_copy</code>	Copy an element
<code>t8_element_compare</code>	Compare the SFC indices of two elements
<code>t8_element_parent</code>	Compute the parent of an element
<code>t8_element_sibling</code>	Compute the $i$ -th sibling (child of parent)
<code>t8_element_num_corners</code>	Return the number of vertices of an element
<code>t8_element_num_faces</code>	Return the number of faces of an element
<code>t8_element_max_num_faces</code>	Return the maximum number of faces of any descendant of an element

<code>t8_element_num_children</code>	Return the number of children of an element
<code>t8_element_num_face_children</code>	Return the number of children at a given face
<code>t8_element_get_face_corner</code>	Given a face $f$ and an index $i$ return the element local index of the $i$ -th corner of $f$
<code>t8_element_get_corner_face</code>	Given a corner $i$ and an index $j$ return the $j$ -th face sharing corner $i$
<code>t8_element_child</code>	Construct the $i$ -th child of an element
<code>t8_element_children</code>	Construct all children of an element
<code>t8_element_child_id</code>	Return the child-id of an element
<code>t8_element_ancestor_id</code>	Return the child-id in the level $\ell$ ancestor
<code>t8_element_is_family</code>	Given a collection of elements, query whether they form a family (that is, they form the output of <code>t8_element_children</code> for some element)
<code>t8_element_nca</code>	Construct the nearest common ancestor of two elements
<code>t8_element_face_class</code>	Return the element type of the $i$ -th face
<code>t8_element_children_at_face</code>	Construct all children that share the $i$ -th face
<code>t8_element_face_child_face</code>	Given a face of an element and the $i$ -th child of that face, return the face number of the child of the element that matches the child face.
<code>t8_element_face_parent_face</code>	Given a face of an element, return the face number of the parent of the element that matches the element's face
<code>t8_element_tree_face</code>	If a given face lies on the tree boundary, return the face number of the respective tree face
<code>t8_element_transform_face</code>	Transform the coordinates of a $(d - 1)$ -dimensional element across a $d$ -dimensional tree-to-tree face connection

<code>t8_element_extrude_face</code>	From a $(d - 1)$ -dimensional element at a $d$ -dimensional tree boundary, construct the respective $d$ -dimensional element within the tree.
<code>t8_element_boundary_face</code>	From a $d$ -dimensional element and a face $f$ construct the $(d - 1)$ -dimensional element representing the face
<code>t8_element_first_descendant_face</code>	Construct the first descendant at a given face
<code>t8_element_last_descendant_face</code>	Construct the last descendant at a given face
<code>t8_element_is_root_boundary</code>	Query whether an element lies at its root tree's boundary across a given face
<code>t8_element_face_neighbor_inside</code>	Construct the face-neighbor of an element within the same tree
<code>t8_element_set_linear_id</code>	Construct an element given a SFC index
<code>t8_element_get_linear_id</code>	Return the SFC index of a given element
<code>t8_element_first_descendant</code>	Construct the first descendant of an element
<code>t8_element_last_descendant</code>	Construct the last descendant of an element
<code>t8_element_successor</code>	Construct the successor (w.r.t. the SFC index) of a given element
<code>t8_element_anchor</code>	Return the (integer) coordinates of the anchor node of an element
<code>t8_element_root_len</code>	Return the (integer) length of the root tree of an element
<code>t8_element_vertex_coords</code>	Return the (integer) coordinates of a given vertex of an element

Table A.1.: All low-level functions provided by `t8code` with a brief description.

# Bibliography

- [1] Mark Ainsworth and Katia Pinchedez. hp-approximation theory for bdfm and rt finite elements on quadrilaterals. *SIAM Journal on Numerical Analysis*, 40(6):2047–2068, 2003.
- [2] Volkan Akçelik, Jacobo Bielak, George Biros, Ioannis Epanomeritakis, Antonio Fernandez, Omar Ghattas, Eui J. Kim, Julio Lopez, David R. O’Hallaron, Tiankai Tu, and John Urbanic. High resolution forward and inverse earthquake modeling on terascale computers. In *SC03: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2003.
- [3] Clelia Albrecht. Parallelization of Adaptive Gradient Augmented Level Set Methods. Master’s thesis, University of Bonn, 2016.
- [4] AMReX. <https://ccse.lbl.gov/AMReX/index.html>.
- [5] Axel Arnold, Olaf Lenz, Stefan Kesselheim, Rudolf Weeber, Florian Fahrenberger, Dominic Roehm, Peter Košován, and Christian Holm. ESPResSo 3.1: Molecular dynamics software for coarse-grained models. In M. Griebel and M. A. Schweitzer, editors, *Meshfree Methods for Partial Differential Equations VI*, volume 89 of *Lecture Notes in Computational Science and Engineering*, pages 1–23. Springer, September 2012.
- [6] I. Babuska and W. C. Rheinboldt. Error estimates for adaptive finite element computations. *SIAM Journal on Numerical Analysis*, 15(4):736–754, 1978.
- [7] M. Bader, H.J. Bungartz, and T. Weinzierl. *Advanced Computing*. Lecture Notes in Computational Science and Engineering. Springer Berlin Heidelberg, 2013.
- [8] M. Bader and Ch. Zenger. *Efficient Storage and Processing of Adaptive Triangular Grids Using Sierpinski Curves*, pages 673–680. Springer, 2006.
- [9] Michael Bader. *Space-Filling Curves: An Introduction with Applications in Scientific Computing*. Texts in Computational Science and Engineering. Springer, 2012.
- [10] Wolfgang Bangerth, Ralf Hartmann, and Guido Kanschat. deal.II – a general-purpose object-oriented finite element library. *ACM Transactions on Mathematical Software*, 33(4):24, 2007.

- [11] Randolph E. Bank, Andrew H. Sherman, and Alan Weiser. Some refinement algorithms and data structures for regular local mesh refinement. In Robert S. Stapleman, editor, *Scientific Computing, Applications of Mathematics and Computing to the Physical Sciences*, volume 1. IMACS/North-Holland Publishing Co., 1983.
- [12] George Keith Batchelor. *An introduction to fluid dynamics*. Cambridge university press, 2000.
- [13] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84, May 1989.
- [14] Marsha J. Berger and Joseph Olinger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53(3):484–512, 1984.
- [15] M. Berzins, J. Luitjens, Q. Meng, T. Harman, and C. A. Wight. Uintah: a scalable framework for hazard analysis. In *Proceedings of the 2010 TeraGrid Conference*, 2010.
- [16] Jürgen Bey. Der BPX-Vorkonditionierer in drei Dimensionen: Gitterverfeinerung, Parallelisierung und Simulation. *Universität Heidelberg*, 1992. Preprint.
- [17] Jürgen Bey. Simplicial grid refinement: on Freudenthal’s algorithm and the optimal number of congruence classes. *Numerische Mathematik*, 85(1):1–29, 2000.
- [18] Boxlib. <https://ccse.lbl.gov/BoxLib/index.html>.
- [19] D. Braess and R. Verfürth. A posteriori error estimators for the raviart-thomas element. *SIAM Journal on Numerical Analysis*, 33(6):2431–2444, 1996.
- [20] Dietrich Braess. *Finite Elements. Theory, Fast Solvers, and Applications in Solid Mechanics*. Cambridge University Press, Cambridge, New York, 1997.
- [21] Kolja Brix, Ralf Massjung, and Alexander Voss. Refinement and connectivity algorithms for adaptive discontinuous Galerkin methods. *SIAM Journal on Scientific Computing*, 33(1):66–101, 2011.
- [22] 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, B. Smith, R. P. Harkness, J. Bordner, J.-h. Kim, M. Kuhlen, H. Xu, N. Goldbaum, C. Hummels, A. G. Kritsuk, E. Tasker, S. Skory, C. M. Simpson, O. Hahn, J. S. Oishi, G. C. So, F. Zhao, R. Cen, Y. Li, and The Enzo Collaboration. ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *ApJS*, 211:19, April 2014.

- [23] Hans-Joachim Bungartz, M. Mehl, and T. Weinzierl. A parallel adaptive Cartesian PDE solver using space-filling curves. *Euro-Par 2006 Parallel Processing*, pages 1064–1074, 2006.
- [24] A. Burri, A. Dedner, R. Klöfkorn, and M. Ohlberger. *An efficient implementation of an adaptive and parallel grid in DUNE*, pages 67–82. Springer, 2006.
- [25] Carsten Burstedde. **p4est**: Parallel AMR on forests of octrees, last accessed November 30, 2017. <http://www.p4est.org/>.
- [26] Carsten Burstedde, Martin Burtscher, Omar Ghattas, Georg Stadler, Tiankai Tu, and Lucas C. Wilcox. ALPS: A framework for parallel adaptive PDE solution. *Journal of Physics: Conference Series*, 180:012009, 2009.
- [27] Carsten Burstedde, Omar Ghattas, Michael Gurnis, Tobin Isaac, Georg Stadler, Tim Warburton, and Lucas C. Wilcox. Extreme-scale AMR. In *SC10: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM/IEEE, 2010.
- [28] Carsten Burstedde, Omar Ghattas, Georg Stadler, Tiankai Tu, and Lucas C. Wilcox. Towards adaptive mesh PDE simulations on petascale computers. In *Proceedings of Teragrid '08*, 2008.
- [29] Carsten Burstedde and Johannes Holke. A tetrahedral space-filling curve for nonconforming adaptive meshes. *SIAM Journal on Scientific Computing*, 38(5):C471–C503, 2016.
- [30] Carsten Burstedde and Johannes Holke. **p4est**: Scalable algorithms for parallel adaptive mesh refinement. In Dirk Brömmel, Wolfgang Frings, and Brian J. N. Wylie, editors, *JUQUEEN Extreme Scaling Workshop 2016*, number FZJ-JSC-IB-2016-01 in JSC Internal Report, pages 49–54. Jülich Supercomputing Centre, 2016.
- [31] Carsten Burstedde and Johannes Holke. Coarse mesh partitioning for tree-based AMR. *SIAM Journal on Scientific Computing*, 39(5):C364–C392, 2017.
- [32] Carsten Burstedde, Johannes Holke, and Tobin Isaac. Bounds on the number of discontinuities of Morton-type space-filling curves. Submitted. Preprint available at <http://arxiv.org/abs/1505.05055>, 2017.
- [33] Carsten Burstedde, Georg Stadler, Laura Alisic, Lucas C. Wilcox, Eh Tan, Michael Gurnis, and Omar Ghattas. Large-scale adaptive mantle convection simulation. *Geophysical Journal International*, 192(3):889–906, 2013.
- [34] Carsten Burstedde, Lucas C. Wilcox, and Omar Ghattas. **p4est**: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011.

- [35] Paul M. Campbell, Karen D. Devine, Joseph E. Flaherty, Luis G. Gervasio, and James D. Teresco. Dynamic octree load balancing using space-filling curves. Technical Report CS-03-01, Williams College Department of Computer Science, 2003.
- [36] C. Carstensen. An adaptive mesh-refining algorithm allowing for an  $H^1$  stable  $L^2$  projection onto Courant finite element spaces. *Constr. Approx.*, 20(4):549–564, 2004.
- [37] U.V. Catalyurek, E.G. Boman, K.D. Devine, D. Bozdag, R.T. Heaphy, and L.A. Riesen. Hypergraph-based dynamic load balancing for adaptive scientific computations. In *Proc. of 21st International Parallel and Distributed Processing Symposium (IPDPS'07)*. IEEE, 2007.
- [38] Henry Ker-Chang Chang and Jiang-Long Liu. A linear quadtree compression scheme for image encryption. *Signal Processing: Image Communication*, 10(4):279–290, 1997.
- [39] C. Chevalier and F. Pellegrini. PT-Scotch: A tool for efficient parallel graph ordering. *Parallel Computing*, 34(6-8):318–331, July 2008.
- [40] Alexandre Joel Chorin. Numerical solution of the navier-stokes equations. *Mathematics of computation*, 22(104):745–762, 1968.
- [41] Albert Cohen, Sidi Mahmoud Kaber, Siegfried Müller, and Marie Postel. Fully adaptive multiresolution finite volume schemes for conservation laws. *Math. Comput.*, 72(241):183–225, January 2003.
- [42] Phillip Colella, Daniel T. Graves, Noel Keen, Terry J. Ligocki, Daniel F. Martin, Peter W. McCorquodale, David Modiano, Peter O. Schwartz, Theodore D. Sternberg, and Brian Van Straalen. *Chombo Software Package for AMR Applications. Design Document*. Applied Numerical Algorithms Group, NERSC Division, Lawrence Berkeley National Laboratory, Berkeley, CA, May 2007.
- [43] Richard Courant, Kurt Friedrichs, and Hans Lewy. Über die partiellen differenzgleichungen der mathematischen physik. *Mathematische annalen*, 100(1):32–74, 1928.
- [44] Richard Courant, Eugene Isaacson, and Mina Rees. On the solution of nonlinear hyperbolic differential equations by finite differences. *Communications on Pure and Applied Mathematics*, 5(3):243–255, 1952.
- [45] Gianluca De Santis. *Novel mesh generation method for accurate image-based computational modelling of blood vessels*. PhD thesis, Ghent University, 2011.
- [46] Karen Devine, Erik Boman, Robert Heaphy, Bruce Hendrickson, and Courtenay Vaughan. Zoltan data management services for parallel dynamic applications. *Computing in Science and Engineering*, 4(2):90–97, 2002.



- [47] Karen D. Devine, Erik G. Boman, Robert T. Heaphy, Bruce A. Hendrickson, James D. Teresco, Jamal Faik, Joseph E. Flaherty, and Luis G. Gervasio. New challenges in dynamic load balancing. *Appl. Numer. Math.*, 52(2–3):133–152, 2005.
- [48] Stefan Donath, Klaus Mecke, Swapna Rabha, Vivek Buwa, and Ulrich Rude. Verification of surface tension in the parallel free surface lattice Boltzmann method in waLBerla. *Computers & Fluids*, 45(1):177–186, 2011. 22nd International Conference on Parallel Computational Fluid Dynamics (ParCFD 2010) ParCFD.
- [49] Willy Dörfler. A convergent adaptive algorithm for Poisson’s equation. *SIAM Journal on Numerical Analysis*, 33(3):1106–1124, 1996.
- [50] Florence Drui, Alexandru Fikl, Pierre Kestener, Samuel Kokh, Adam Larat, Vincent Le Chenadec, and Marc Massot. Experimenting with the p4est library for amr simulations of two-phase flows. *ESAIM: Proceedings and Surveys*, 53:232–247, 2016.
- [51] enGrid. <https://github.com/enGits/engrid>.
- [52] Daming Feng, Christos Tsolakis, Andrey N Chernikov, and Nikos P Chrisochoides. Scalable 3d hybrid parallel delaunay image-to-mesh conversion algorithm for distributed shared memory architectures. *Computer-Aided Design*, 85:10–19, 2017.
- [53] Paul F. Fischer, Gerald W. Kruse, and Francis Loth. Spectral element methods for transitional flows in complex geometries. *Journal of Scientific Computing*, 17(1-4):81–98, December 2002.
- [54] Panagiotis Foteinos and Nikos Chrisochoides. 4d space—time delaunay meshing for medical images. *Eng. with Comput.*, 31(3):499–511, July 2015.
- [55] R. A. Gentry, R. E. Martin, and B. J. Daly. An eulerian differencing method for unsteady compressible flow problems. *Journal of Computational Physics*, 1:87–118, August 1966.
- [56] Kai Germaschewski, Amitava Bhattacharjee, Rainer Grauer, and Barry Smith. Using krylov-schwarz methods in an adaptive mesh refinement environment. In *Adaptive Mesh Refinement - Theory and Applications*, pages 115–124. Springer Science & Business Media, 2003.
- [57] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.
- [58] N.A. Golias and R.W. Dutton. Delaunay triangulation and 3D adaptive mesh generation. *Finite Elements in Analysis and Design*, 25(3-4):331–341, 1997. Adaptive Meshing, Part 2.

- [59] M. Griebel and G. Zumbusch. Parallel multigrid in an adaptive PDE solver based on hashing. In E. H. D'Hollander, G. R. Joubert, F. J. Peters, and U. Trottenberg, editors, *Parallel Computing: Fundamentals, Applications and New Directions, Proceedings of the Conference ParCo'97, 19–22 September 1997, Bonn, Germany*, volume 12, pages 589–600. Elsevier, North-Holland, 1998.
- [60] Michael Griebel and Gerhard W. Zumbusch. Parallel multigrid in an adaptive PDE solver based on hashing and space-filling curves. *Parallel Computing*, 25:827–843, 1999.
- [61] Arthur Guittet, Maxime Theillard, and Frédéric Gibou. A stable projection method for the incompressible navier-stokes equations on arbitrary geometries and adaptive quad/octrees. *Journal of Computational Physics*, 292:215 – 238, 2015.
- [62] Frank Günther, Miriam Mehl, Markus Pögl, and Christoph Zenger. A cache-aware algorithm for PDEs on hierarchical data structures based on space-filling curves. *SIAM Journal on Scientific Computing*, 28(5):1634–1650, 2006.
- [63] Herman Haverkort. Sixteen space-filling curves and traversals for d-dimensional cubes and simplices. *arXiv preprint arXiv:1711.04473*, 2017.
- [64] Herman Haverkort and Freek van Walderveen. Locality and bounding-box quality of two-dimensional space-filling curves. *Computational Geometry*, 43(2):131–174, 2010.
- [65] D. Hilbert. Über die stetige Abbildung einer Linie auf ein Flächenstück. *Mathematische Annalen*, 38:459–460, 1891.
- [66] Ø. Hjelle and M. Dæhlen. *Triangulations and Applications*. Mathematics and Visualization. Springer, 2006.
- [67] T.J.R. Hughes, J.A. Cottrell, and Y. Bazilevs. Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement. *CMAME*, 2005.
- [68] Daniel A. Ibanez, E. Seogyong Seol, Cameron W. Smith, and Mark S. Shephard. Pumi: Parallel unstructured mesh infrastructure. *ACM Trans. Math. Softw.*, 42(3):17:1–17:28, May 2016.
- [69] Tobin Isaac, Carsten Burstedde, and Omar Ghattas. Low-cost parallel algorithms for 2:1 octree balance. In *Proceedings of the 26th IEEE International Parallel & Distributed Processing Symposium*. IEEE, 2012. <http://dx.doi.org/10.1109/IPDPS.2012.47>.
- [70] Tobin Isaac, Carsten Burstedde, Lucas C. Wilcox, and Omar Ghattas. Recursive algorithms for distributed forests of octrees. *SIAM Journal on Scientific Computing*, 37(5):C497–C531, 2015.

- [71] Yasushi Ito, Alan M Shih, Anil K Erukala, Bharat K Soni, Andrey Chernikov, Nikos P Chrisochoides, and Kazuhiro Nakahashi. Parallel unstructured mesh generation by an advancing front method. *Mathematics and Computers in Simulation*, 75(5):200–209, 2007.
- [72] Claes Johnson. Error estimates and adaptive time-step control for a class of one-step methods for stiff ordinary differential equations. *SIAM Journal on Numerical Analysis*, 25(4):908–926, 1988.
- [73] Jülich Supercomputing Centre. JUQUEEN: IBM Blue Gene/Q supercomputer system at the Jülich Supercomputing Centre. *Journal of large-scale research facilities*, A1, 2015.
- [74] George Karypis and Vipin Kumar. *METIS – Unstructured Graph Partitioning and Sparse Matrix Ordering System, Version 2.0*, 1995.
- [75] George Karypis and Vipin Kumar. A parallel algorithm for multilevel graph partitioning and sparse matrix ordering. *Journal of Parallel and Distributed Computing*, 48:71–95, 1998.
- [76] Aly Khawaja and Yannis Kallinderis. Hybrid grid generation for turbomachinery and aerospace applications. *International Journal for Numerical Methods in Engineering*, 49(1-2):145–166, 2000.
- [77] Andrew C Kirby, Michael J Brazell, Zhi Yang, Rajib Roy, Behzad R Ahrabi, Dimitri J Mavriplis, Michael K Stoellinger, and Jay Sitaraman. Wind farm simulations using an overset hp-adaptive approach with blade-resolved turbine models. In *23rd AIAA Computational Fluid Dynamics Conference*, page 3958, 2017.
- [78] Benjamin S. Kirk, John W. Peterson, Roy H. Stogner, and Graham F. Carey. `libMesh`: A C++ library for parallel adaptive mesh refinement/coarsening simulations. *Engineering with Computers*, 22(3–4):237–254, 2006.
- [79] Richard I. Klein. Star formation with 3-d adaptive mesh refinement: the collapse and fragmentation of molecular clouds. *Journal of Computational and Applied Mathematics*, 109(1):123 – 152, 1999.
- [80] M. Klitz. *Numerical Simulation of Droplets with Dynamic Contact Angles*. PhD thesis, Institut für Numerische Simulation, University of Bonn, December 2014.
- [81] David Knapp. Adaptive Verfeinerung von Prismen. Bachelor’s thesis, University of Bonn, 2017.
- [82] David A. Kopriva, Stephen L. Woodruff, and M. Yousuff Hussaini. Computation of electromagnetic scattering with a non-conforming discontinuous spectral element method. *International Journal for Numerical Methods in Engineering*, 53(1):105–122, 2002.

- [83] F. Koster, M. Griebel, N. Kevlahan, M. Farge, and K. Schneider. Towards an adaptive wavelet-based 3D Navier-Stokes solver. In E.H. Hirschel, editor, *Numerical flow simulation I, Notes on Numerical Fluid Mechanics, Vol. 66*, pages 339–364. Vieweg-Verlag, Braunschweig, 1998.
- [84] Martin Kronbichler, Timo Heister, and Wolfgang Bangerth. High accuracy mantle convection simulation through modern numerical methods. *Geophysical Journal International*, 191(1):12–29, 2012.
- [85] Michael Lahnert, Carsten Burstedde, Christian Holm, Miriam Mehl, Georg Rempfer, and Florian Weik. Towards lattice-Boltzmann on dynamically adaptive grids—minimally-invasive grid exchange in ESPResSo. In M. Papadrakakis, V. Papadopoulos, G. Stefanou, and V. Plevris, editors, *ECCOMAS Congress 2016, VII European Congress on Computational Methods in Applied Sciences and Engineering*, pages 1–25. ECCOMAS, 2016.
- [86] Andras Laszloffy, Jingping Long, and Abani K. Patra. Simple data management, scheduling and solution strategies for managing the irregularities in parallel adaptive hp finite element simulations. *Parallel Computing*, 26:1765–1788, 2000.
- [87] Orion S. Lawlor, Sayantan Chakravorty, Terry L. Wilmarth, Niles Choudhury, Isaac Dooley, Gengbin Zheng, and Laxmikant V. Kalé. ParFUM: a parallel framework for unstructured meshes for scalable dynamic physics applications. *Engineering with Computers*, 22(3):215–235, 2006.
- [88] Henri Léon Lebesgue. *Leçons sur l'intégration et la recherche des fonctions primitives*. Gauthier-Villars, 1904.
- [89] M. Lee, L. De Floriani, and H. Samet. Constant-time neighbor finding in hierarchical tetrahedral meshes. In *SMI 2001 International Conference on Shape Modeling and Applications*, pages 286–295, May 2001.
- [90] Randall J. LeVeque. *Finite volume methods for hyperbolic problems*. Cambridge University Press, 2002.
- [91] Z. Li and K. Ito. *The Immersed Interface Method*. Society for Industrial and Applied Mathematics, 2006.
- [92] Anwei Liu and Barry Joe. Quality local refinement of tetrahedral meshes based on 8-subtetrahedron subdivision. *Math. Comput.*, 65(215):1183–1200, July 1996.
- [93] Anders Logg, Kent-Andre Mardal, and Garth N. Wells, editors. *Automated Solution of Differential Equations by the Finite Element Method*, volume 84 of *Lecture Notes in Computational Science and Engineering*. Springer, 2012.

- [94] David Marcum. Generation of unstructured grids for viscous flow applications. In *33rd Aerospace Sciences Meeting and Exhibit*, page 212, 1995.
- [95] DG Martineau, Sean Stokes, SJ Munday, AP Jackson, BJ Gribben, and Niek Verhoeven. Anisotropic hybrid mesh generation for industrial rans applications. *AIAA Paper*, 534(2006):39, 2006.
- [96] Oliver Meister, Kaveh Rahnema, and Michael Bader. A software concept for cache-efficient simulation on dynamically adaptive structured triangular grids. In *PARCO*, pages 251–260, 2011.
- [97] Oliver Meister, Kaveh Rahnema, and Michael Bader. Parallel memory-efficient adaptive mesh refinement on structured triangular meshes with billions of grid cells. *ACM Trans. Math. Softw.*, 43(3):19:1–19:27, September 2016.
- [98] Chohong Min. On reinitializing level set functions. *Journal of Computational Physics*, 229(8):2764 – 2772, 2010.
- [99] Mohammad Mirzadeh, Arthur Guittet, Carsten Burstedde, and Frédéric Gibou. Parallel level-set methods on adaptive tree-based grids. *Journal of Computational Physics*, 322:345–364, 2016.
- [100] William F. Mitchell. A refinement-tree based partitioning method for dynamic load balancing with adaptively refined grids. *Journal of Parallel and Distributed Computing*, 67(4):417–429, 2007.
- [101] G. M. Morton. A computer oriented geodetic data base; and a new technique in file sequencing. Technical report, IBM Ltd., 1966.
- [102] Andreas Müller, Jörn Behrens, Francis X. Giraldo, and Volkmar Wirth. Comparison between adaptive and uniform discontinuous galerkin simulations in dry 2d bubble experiments. *Journal of Computational Physics*, 235:371 – 393, 2013.
- [103] Andreas Müller, Michal A. Kopera, Simone Marras, Lucas C. Wilcox, Tobin Isaac, and Francis X. Giraldo. Strong scaling for numerical weather prediction at petascale with the atmospheric model NUMA. <http://arxiv.org/abs/1511.01561>, 2015.
- [104] Siegfried Müller and Youssef Stiriba. Fully adaptive multiscale schemes for conservation laws employing locally varying time stepping. *Journal of Scientific Computing*, 30(3):493–531, 2007.
- [105] J. Müller, O. Sahni, X. Li, K. E. Jansen, M. S. Shephard, and C. A. Taylor. Anisotropic adaptive finite element method for modelling blood flow. *Computer Methods in Biomechanics and Biomedical Engineering*, 8(5):295–305, 2005. PMID: 16298851.

- [106] Charles D. Norton, Greg Lyzenga, Jay Parker, and Robert E. Tisdale. Developing parallel GeoFEST(P) using the PYRAMID AMR library. Technical report, Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2004.
- [107] Leonid Oliker and Rupak Biswas. PLUM: Parallel load balancing for adaptive unstructured meshes. *Journal of Parallel and Distributed Computing*, 52(2):150–177, 1998.
- [108] Leonid Oliker, Rupak Biswas, and Harold N. Gabow. Parallel tetrahedral mesh adaptation with dynamic load balancing. *Parallel Computing*, 26(12):1583–1608, November 2000.
- [109] OpenCFD. *OpenFOAM – The Open Source CFD Toolbox – User’s Guide*. OpenCFD Ltd., United Kingdom, 1.4 edition, 11 2007.
- [110] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton Jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [111] Stanley J. Osher and Ronald P. Fedkiw. *Level Set Methods and Dynamic Implicit Surfaces*. Springer, 2002.
- [112] Ekow J. Otoo and Hongwen Zhu. Indexing on spherical surfaces using semi-quadcodes. In David Abel and Beng Chin Ooi, editors, *Advances in Spatial Databases*, volume 692 of *Lecture Notes in Computer Science*, pages 510–529. Springer, 1993.
- [113] Guisepepe Peano. Sur une courbe, qui remplit toute une aire plane. *Math. Ann.*, 36(1):157–160, 1890.
- [114] Daniel Peter, Dimitri Komatitsch, Yang Luo, Roland Martin, Nicolas Le Goff, Emanuele Casarotti, Pieyre Le Loher, Federica Magnoni, Qinya Liu, Celine Blitz, Tarje Nisson-Meyer, Piero Basini, and Jeroen Tromp. Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes. *Geophysical Journal International*, 186(2):721–739, 2011.
- [115] James C. Phillips, Rosemary Braun, Wei Wang, James Gumbart, Emad Tajkhorshid, Elizabeth Villa, Christophe Chipot, Robert D. Skeel, Laxmikant Kalé, and Klaus Schulten. Scalable molecular dynamics with namd. *Journal of Computational Chemistry*, 26(16):1781–1802, 2005.
- [116] John R. Pilkington and Scott B. Baden. Partitioning with spacefilling curves. Technical report, Dept. of Computer Science and Engineering, University of California, San Diego, 1994.
- [117] Ali Pinar and Cevdet Aykanat. Fast optimal load balancing algorithms for 1D partitioning. *Journal on Parallel and Distributed Computing*, 64(8):974–996, August 2004.

- [118] Frank Ramsak, Volker Markl, Robert Fenk, Martin Zirkel, Klaus Elhardt, Rudolf Bayer, Bayerisches Forschungszentrum, and Tu München. Integrating the ub-tree into a database system kernel. In *Proceedings of the 26th International Conference on Very Large Data Bases*, pages 263–272, 2000.
- [119] M. Rasquin, C. Smith, K. Chitale, E. S. Seol, B. A. Matthews, J. L. Martin, O. Sahni, R. M. Loy, M. S. Shephard, and K. E. Jansen. Scalable implicit flow solver for realistic wing simulations with flow control. *Computing in Science Engineering*, 16(6):13–21, Nov 2014.
- [120] Werner C. Rheinboldt and Charles K. Mesztenyi. On a data structure for adaptive finite element mesh refinements. *ACM Transactions on Mathematical Software*, 6(2):166–187, 1980.
- [121] RC Ripley, F-S Lien, and MM Yovanovich. Adaptive unstructured mesh refinement of supersonic channel flows. *International Journal of Computational Fluid Dynamics*, 18(2):189–198, 2004.
- [122] Johann Rudi, A. Cristiano I. Malossi, Tobin Isaac, Georg Stadler, Michael Gurnis, Peter W. J. Staar, Yves Ineichen, Costas Bekas, Alessandro Curioni, and Omar Ghattas. An extreme-scale implicit solver for complex pdes: highly heterogeneous flow in earth’s mantle. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, page 5. ACM, 2015.
- [123] Hans Sagan. *Space-Filling Curves*. Springer, 1994.
- [124] Rahul S. Sampath, Santi S. Adavani, Hari Sundar, Ilya Lashuk, and George Biros. **Dendro**: Parallel algorithms for multigrid and AMR methods on 2:1 balanced octrees. In *SC’08: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2008.
- [125] Joachim Schöberl. Netgen an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and Visualization in Science*, 1(1):41–52, Jul 1997.
- [126] P. M. Selwood and M. Berzins. Parallel unstructured tetrahedral mesh adaptation: algorithms, implementation and scalability. *Concurrency: Practice and Experience*, 11(14):863–884, 1999.
- [127] Jonathan Richard Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. In Ming C. Lin and Dinesh Manocha, editors, *Applied Computational Geometry: Towards Geometric Engineering*, volume 1148 of *Lecture Notes in Computer Science*, pages 203–222. Springer, 1996. From the First ACM Workshop on Applied Computational Geometry.

- [128] Hang Si. *TetGen—A Quality Tetrahedral Mesh Generator and Three-Dimensional Delaunay Triangulator*. Weierstraß Institute for Applied Analysis and Stochastics, Berlin, 2006.
- [129] Waław Sierpiński. Sur une nouvelle courbe continue qui remplit toute une aire plane. *Bulletin de l'Académie des Sciences de Cracovie*, Séries A:462–478, 1912.
- [130] William Skamarock, Joseph Oliger, and Robert L Street. Adaptive grid refinement for numerical weather prediction. *jcp*, 80(1):27 – 60, 1989.
- [131] C. W. Smith, M. Rasquin, D. Ibanez, K. E. Jansen, and M. S. Shephard. Application specific mesh partition improvement. Technical Report 2015-3, Rensselaer Polytechnic Institute, 2015.
- [132] Pavel Šolín, Jakub Červený, and Ivo Doležel. Arbitrary-level hanging nodes and automatic adaptivity in the hp-fem. *Mathematics and Computers in Simulation*, 77(1):117 – 132, 2008.
- [133] David A Steinman, Jaques S Milner, Chris J Norley, Stephen P Lownie, and David W Holdsworth. Image-based computational simulation of flow dynamics in a giant intracranial aneurysm. *American Journal of Neuroradiology*, 24(4):559–566, 2003.
- [134] James R. Stewart and H. Carter Edwards. A framework approach for developing parallel adaptive multiphysics applications. *Finite Elements in Analysis and Design*, 40(12):1599–1617, 2004.
- [135] E. Suli, C. Schwab, and P. Houston. *Discontinuous Galerkin Methods. Theory, Computation and Applications*, chapter hp-DGFEM for partial differential equations with nonnegative characteristic form, pages 221–230. Lecture notes in Computational Science and Engineering. Springer-Verlag, 2000.
- [136] Hari Sundar, Rahul Sampath, and George Biros. Bottom-up construction and 2:1 balance refinement of linear octrees in parallel. *SIAM Journal on Scientific Computing*, 30(5):2675–2708, 2008.
- [137] Hari Sundar, Rahul S. Sampath, Santi S. Adavani, Christos Davatzikos, and George Biros. Low-constant parallel algorithms for finite element simulations using linear octrees. In *SC'07: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2007.
- [138] Mark Sussman, Peter Smereka, and Stanley Osher. A level set approach for computing solutions to incompressible two-phase flow. *Journal of Computational Physics*, 114:146–159, 1994.
- [139] t8code. <https://github.com/holke/t8code>.



- [140] T. J. Tautges, R. Meyers, K. Merkley, C. Stimpson, and C. Ernst. MOAB: A mesh-oriented database. SAND2004-1592, Sandia National Laboratories, April 2004. Report.
- [141] Timothy J Tautges, Jason A Kraftcheck, Nathan Bertram, Vipin Sachdeva, and John Magerlein. Mesh interface resolution and ghost exchange in a parallel mesh representation. In *Parallel and Distributed Processing Symposium Workshops & PhD Forum (IPDPSW), 2012 IEEE 26th International*, pages 1670–1679. IEEE, 2012.
- [142] Herbert Tropf and H. Herzog. Multidimensional range search in dynamically balanced trees. *Angewandte Informatik*, 2:71–77, 1981.
- [143] Tiankai Tu and David R O’Hallaron. Balance refinement of massive linear octree datasets. Technical Report CMU-CS-04, Carnegie Mellon University, Pittsburgh, 2004.
- [144] Tiankai Tu, David R. O’Hallaron, and Omar Ghattas. Scalable parallel octree meshing for terascale applications. In *SC ’05: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2005.
- [145] Vinod Valsalam and Anthony Skjellum. A framework for high-performance matrix multiplication based on hierarchical abstractions, algorithms and optimized low-level kernels. *Concurrency and Computation: Practice and Experience*, 14(10):805–839, 2002.
- [146] David A. Venditti and David L. Darmofal. Adjoint error estimation and grid adaptation for functional outputs: Application to quasi-one-dimensional flow. *Journal of Computational Physics*, 164:204–227, 2000.
- [147] Samir Vinchurkar and P Worth Longest. Evaluation of hexahedral, prismatic and hybrid mesh styles for simulating respiratory aerosol dynamics. *Computers & Fluids*, 37(3):317–331, 2008.
- [148] Timothy Warburton and Thomas Hagstrom. Taming the cfl number for discontinuous galerkin methods on structured meshes. *SIAM Journal on Numerical Analysis*, 46(6):3151–3180, 2008.
- [149] Tobias Weinzierl and Miriam Mehl. Peano—a traversal and storage scheme for octree-like adaptive Cartesian multiscale grids. *SIAM Journal on Scientific Computing*, 33(5):2732–2760, October 2011.
- [150] Lucas C. Wilcox, Georg Stadler, Carsten Burstedde, and Omar Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media. *Journal of Computational Physics*, 229(24):9373–9396, 2010.

- [151] A. M. Wissink, R. D. Hornung, S. R. Kohn, S. S. Smith, and N. S. Ellion. Large scale structured AMR calculations using the SAMRAI framework. In *Proceedings of Supercomputing, ACM/IEEE 2001 Conference*. ACM/IEEE, Nov. 10-16 2001.
- [152] Yusuf Yilmaz, Can Özturan, Oğuz Tosun, Ali Haydar Özer, and Seren Soner. Parallel mesh generation, migration and partitioning for the elmer application. Technical report, PREMA-Partnership for Advanced Computing in Europe, 2010.
- [153] S. Zhang. Successive subdivisions of tetrahedra and multigrid methods on tetrahedral meshes. *Houston Journal of Mathematics*, 21(3), 1995.
- [154] YM Zhang, Y Gu, JT Chen, et al. Boundary layer effect in BEM with high order geometry elements using transformation. *Computer Modeling in Engineering and Sciences (CMES)*, 45(3):227, 2009.
- [155] O. C. Zienkiewicz, D. W. Kelley, J. Gago, and I. Babuska. Hierarchical finite element approaches, error estimates and adaptive refinement. In J. R. Whiteman, editor, *The Mathematics of Finite Elements and Applications IV, MAFELAP 1981*, pages 313–346. Academic Press, London, New York, 1982.
- [156] G. Zumbusch. Dynamic load balancing in a lightweight adaptive parallel multigrid PDE solver. In B. Hendrickson, K. Yelick, C. Bischof, I. Duff, A. Edelman, G. Geist, M. Heath, M. Heroux, C. Koelbel, R. Schrieber, R. Sinovec, and M. Wheeler, editors, *Proceedings of 9th SIAM Conference on Parallel Processing for Scientific Computing (PP '99), San Antonio, Texas*, ISBN 0-89871-435-4, page 10, Philadelphia, PA, 1999. SIAM.
- [157] G. Zumbusch. *Parallel Multilevel Methods. Adaptive Mesh Refinement and Load Balancing*. Teubner, 2003.
- [158] Gerhard W. Zumbusch. On the quality of space-filling curve induced partitions, 2000. *Zeitschrift für Angewandte Mathematik und Mechanik*.
- [159] Gerhard W. Zumbusch. Load balancing for adaptively refined grids. *Proceedings in Applied Mathematics and Mechanics*, 1(1), 2002.

# List of Tables

4.1.	The types of the children of a simplex . . . . .	44
4.2.	The local index of the children of a simplex . . . . .	53
4.3.	Face neighbors for triangles . . . . .	59
4.4.	Face neighbors for tetrahedra . . . . .	59
4.5.	The coordinates $x_i, x_j$ , and $x_k$ for the computation whether a simplex is ancestor of another. . . . .	62
4.6.	The local index of a tetrahedron . . . . .	66
4.7.	The cube-id in dependence of the local index and the type of the parent	66
4.8.	The type in dependence of the local index and the type of the parent	66
5.1.	The relative counts of face-connected and non-connected segments of Morton SFCs . . . . .	81
6.1.	Basic definitions . . . . .	86
6.2.	Runtimes for <code>Partition_cmesh</code> on JUQUEEN . . . . .	109
6.3.	Runtimes for <code>Partition_cmesh</code> with small meshes . . . . .	110
6.4.	Runtimes of coarse mesh repartitioning with 8,192 MPI ranks . . . . .	113
6.5.	Runtimes of forest mesh repartitioning with 8,192 MPI ranks . . . . .	113
6.6.	Runtimes of coarse and forest mesh repartitioning with 458,752 MPI ranks . . . . .	113
7.1.	Face number and type for tetrahedral subfaces . . . . .	121
7.2.	<code>t8_element_boundary_face</code> for quadrilaterals and hexahedra. . . . .	122
7.3.	Coordinates and types for the faces of simplices . . . . .	123
7.4.	$\text{sign}_{t,t'}(g, g')$ from Definition 7.13 for two possible tree-to-tree con- nections. . . . .	125
7.5.	<code>t8_transform_face</code> ( $F, o, s = 1$ ). . . . .	125
7.6.	<code>t8_transform_face</code> ( $F, o = 0, s = -1$ ). . . . .	126
7.7.	The computation of <code>t8_element_extrude_face</code> ( $F', T', g'$ ) . . . . .	128
7.8.	The child indices of all children of an element touching a given face.	129
7.9.	Runtime tests for the three different <code>Ghost</code> algorithms. . . . .	141
8.1.	Strong scaling with tetrahedral elements. . . . .	149
8.2.	Weak scaling with tetrahedral and hexahedral elements. . . . .	151
9.1.	Verifying convergence of the advection solver . . . . .	167
9.2.	Parallel efficiency of the advection solver . . . . .	170
9.3.	Comparing uniform and adaptive meshes. . . . .	171

A.1. All low-level functions provided by `t8code` with a brief description. . 182

## List of Figures

2.1.	Uniform and adaptive mesh . . . . .	17
2.2.	Three different AMR methods . . . . .	18
2.3.	A refined quadrilateral and the associated refinement tree. . . . .	20
2.4.	Two ways to approximate a non-trivial domain. . . . .	21
2.5.	AMR algorithms pipeline . . . . .	24
3.1.	1:4 refinement of a quadrilateral element. . . . .	29
3.2.	Child-ids for the cubical Morton index . . . . .	31
3.3.	Computing the Morton index of a quad. . . . .	32
3.4.	The SFC curve arising from the Morton code for quadrilateral 1:4 refinement . . . . .	33
3.5.	Complex geometries with multiple trees . . . . .	35
4.1.	The Bey refinement schemes for triangles and tetrahedra . . . . .	38
4.2.	The triangle and tetrahedra types . . . . .	40
4.3.	First refining and then triangulating a cube is the same as first tri- angulating it and then refining the tetrahedra . . . . .	41
4.4.	Resolving hanging nodes . . . . .	42
4.5.	The TM curve preserves locality . . . . .	47
4.6.	The cube-id . . . . .	48
4.7.	A recursive description of the TM curve . . . . .	54
4.8.	The type of the parent of a simplex . . . . .	56
4.9.	The face-neighbors of a tetrahedron . . . . .	58
4.10.	Determining when a triangle is outside of the root triangle . . . . .	63
4.11.	Runtime tests for <code>New</code> on <code>JUQUEEN</code> . . . . .	72
4.12.	Strong scaling for <code>Adapt</code> on <code>JUQUEEN</code> . . . . .	72
5.1.	Two example segments of the TM curve . . . . .	76
5.2.	Illustration of Lemma 5.4 . . . . .	77
5.3.	The 2D case in the proof of Proposition 5.8 . . . . .	79
5.4.	The relative count of TM-curve segments by number of face-connected components . . . . .	82
6.1.	Coarse and fine mesh . . . . .	85
6.2.	Complex geometries with multiple trees . . . . .	86
6.3.	Vertex and face labels of the different tree types . . . . .	87
6.4.	A partitioned coarse mesh with shared trees . . . . .	90
6.5.	An example for coarse mesh repartitioning . . . . .	97

6.6.	The communication patterns for the different types of face information	101
6.7.	The coarse mesh of disjoint bricks . . . . .	107
6.8.	Weak scaling of coarse mesh repartitioning . . . . .	108
6.9.	Strong scaling of coarse mesh repartitioning . . . . .	108
6.10.	A coarse mesh that models a brick with holes . . . . .	112
6.11.	The refined forest on the brick with holes mesh . . . . .	112
7.1.	Element face-neighbor . . . . .	118
7.2.	Two hexahedral elements that are face-neighbors across tree boundaries. . . . .	119
7.3.	Constructing the face element $F$ to an element $E$ at a tree face $G$ . . . . .	121
7.4.	An element $E$ and its leaf elements. . . . .	131
7.5.	An example for <code>t8_owners_at_face</code> . . . . .	132
7.6.	Comparing the different implementations of <code>Ghost</code> . . . . .	140
8.1.	An example for face-balance with a triangular forest and 5 MPI ranks.	143
8.2.	Testing <code>Ghost</code> and <code>Ripple-balance</code> on a unit cube geometry. . . . .	148
8.3.	Strong scaling with tetrahedral elements. . . . .	150
8.4.	Weak scaling with tetrahedral and hexahedral elements. . . . .	152
9.1.	Illustration for flux computations. . . . .	157
9.2.	2D unit square for advection (quad/triangle/hybrid). . . . .	162
9.3.	Simulation of the 2D test case on a 2D hybrid mesh. . . . .	163
9.4.	Illustration of a 3D flow. . . . .	164
9.5.	A cutout view of a hybrid 3D mesh . . . . .	165
9.6.	Strong scaling results for the advection solver . . . . .	168
9.7.	Strong scaling results for the advection solver (full JUQUEEN) . . . . .	169
9.8.	2D flow around a disk. . . . .	172
9.9.	The coarse mesh that we use to model the domain with a disk cut out.	172
9.10.	Refinement criterion with different coarse mesh element sizes. . . . .	173
9.11.	Solutions to the advection around a disk and the adapted mesh. . . . .	174