

POINTERLESS IMPLEMENTATION OF HIERARCHICAL SIMPLICIAL MESHES AND EFFICIENT NEIGHBOR FINDING IN ARBITRARY DIMENSIONS*

F. Betul Atalay

David M. Mount

*Department of Computer Science, University of Maryland, College Park, MD 20742.
{betul,mount}@cs.umd.edu*

ABSTRACT

We describe a pointerless representation of hierarchical regular simplicial meshes, based on a bisection approach proposed by Maubach. We introduce a new labeling scheme, called an LPT code, which uniquely encodes each simplex of the hierarchy. We present rules to compute the neighbors of a given simplex efficiently through the use of these codes. In addition, we show how to traverse the associated tree and how to answer point location and interpolation queries. Our system works in arbitrary dimensions.

Keywords: Pointerless data structures, neighbor finding, hierarchical meshes, simplicial meshes, interpolation.

1. INTRODUCTION

Hierarchical simplicial meshes have been widely used in various application areas such as finite element computations, scientific visualization and geometric modeling. There has been considerable amount of work in simplicial mesh refinement, particularly in 2- and 3-dimensions, and a number of different refinement techniques have been proposed [1, 2, 3, 4, 5, 6, 7]. Because of the need to handle data sets with temporal components, there is a growing interest in higher dimensional meshes. In this paper, we build on a bisection refinement method proposed by Maubach [6].

A hierarchical mesh is said to be *regular* if the vertices of the mesh are regularly distributed and the process by which a cell is subdivided is identical for all cells. Maubach developed a simple bisection algorithm based on a particular ordering of vertices and presented a mathematically rigorous analysis of the geometric structure of the hierarchical regular simplicial

meshes in any dimension d [6]. The mesh is generated by a process of repeated bisection applied to a hypercube that has been initially subdivided into $d!$ congruent simplices. The subdivision pattern repeats itself on a smaller scale at every d levels. Whenever a simplex is bisected, some of its neighboring simplices may need to be bisected as well, in order to guarantee that the entire subdivision is *compatible*. Intuitively, a *compatible* subdivision is a subdivision in which pairs of neighboring cells meet along a single common face. A compatible simplicial subdivision is also referred to as a *simplicial complex* [8].

In computer graphics, adaptively refined regular meshes in 2- or 3-dimensions have been of interest for their use in realistic surface and volume rendering [9, 10]. In many such applications, efficiency of various operations such as traversal and neighbor finding on the mesh is most desired. Based on the 3-dimensional version of Maubach's method, Hebert [11] presented a more efficient symbolic implementation of regular tetrahedral meshes by introducing an addressing scheme that allows unique labeling of the tetrahedra in the mesh, and he showed how to compute

*THIS MATERIAL IS BASED UPON WORK SUPPORTED BY THE NATIONAL SCIENCE FOUNDATION UNDER GRANT NO. 0098151.

the face neighbors of a tetrahedron based on its label. Hebert’s addressing scheme could be generalized to higher dimensions, however the neighbor finding algorithms are quite specific to 3-dimensions, and a generalization to higher dimensions is a definite challenge. In this paper, we present such an algorithm that works in arbitrary dimensions.

Our interest in higher dimensions is motivated by another computer graphics application that accelerates ray-tracing [12] through multi-dimensional interpolation. In this application, rays in 3-space are modeled as points in a 4-dimensional parameter space, and each sample ray is traced through a scene to gather various geometric attributes that are required to compute an intensity value. Since tracing a ray through a complex scene can be computationally intensive, our approach is to instead sample a relatively small number of rays, and then interpolate among these samples to reconstruct the value at intermediate rays [13, 14]. An adaptively refined regular simplicial mesh is constructed over the 4-dimensional domain of interest, and the field values are sampled at the vertices of this subdivision. For interpolation purposes, compatibly refined simplicial meshes are preferable over octree subdivisions, since they guarantee continuous interpolants, and that they are much simpler in the sense that interpolations are performed with a minimal number of samples.

In addition to our own motivation, higher dimensional meshes are of interest for visualization of time-varying fields, and efficient algorithms for performing traversals and neighbor finding is required.

Thus, our main objective in this paper is to present an efficient implementation of hierarchical regular simplicial meshes in any dimension d . Rather than representing the hierarchy explicitly as a tree using child pointers, we use a *pointerless representation* in which we access nodes through an index called a *location code*. Location codes [15, 16] have arisen as a popular alternative to standard pointer-based representations, because they separate the hierarchy from its representation, and so allow the application of very efficient access methods. The space savings realized by not having to store pointers (to the parent, two children, and $d + 1$ neighboring simplices) and simplex vertices is quite significant for large multidimensional meshes.

We present a location code, called the *LPT code*, which can be used to access nodes of this tree. We store the mesh in a data structure called a *simplex decomposition tree*. Our hierarchical decomposition is based on the same bisection method given by Maubach [6]. (Note that Maubach’s representation is not pointerless.) In addition to efficient computation of neighbors, we show how to perform tree traversals, point locations, and answer interpolation queries efficiently through the use of these codes.

2. POINTERLESS REPRESENTATIONS AND PRIOR WORK

Regular subdivisions have the disadvantage of limiting the mesh’s ability to adapt to the variational structure of the scalar field, but they provide a number of significant advantages from the perspectives of efficiency, practicality, and ease of use. These include guarantees on geometric quality, relief from the need to explicitly storing topological information, straightforward methods for performing point location and other queries. One very practical advantage of regularity involves performance issues arising from modern memory hierarchies, which are based on multiple levels ranging from registers and caches to main memory and disk (including virtual memory). The storage capacity at each level increases, and so too does access latency. There are often many orders of magnitude of difference between the time needed to access local data (which may be stored in registers or cache) versus global data (which may reside on disk) [17]. Large dynamic pointer-based data structures are particularly problematic from this perspective, because node storage is typically allocated and deallocated dynamically and, unless special care is taken, simple pointer-based traversals suffer from a nonlocal pattern of memory references. This is one of the principal motivating factors behind I/O efficient algorithms [18] and cache conscious data structures [17].

In contrast with pointer-based implementations, regular spatial subdivisions support *pointerless* implementations. Pointerless versions of quadtree and its variants have been known for many years [19, 15]. The idea is to associate each node of the tree with a unique index, called a *location code*. Because of the regularity of the subdivision, given any point in space, it is possible to compute the location code of the node of a particular depth in the tree that contains this point. This can be done entirely in local memory, without accessing the data structure in global memory. Once the location code is known, the actual node can be accessed through a small number of accesses to global memory (e.g., by hashing).

Prior work on pointerless regular simplicial meshes has principally been in 2- and 3-dimensions. Lee and Samet presented a pointerless hierarchical triangulation based on a four-way decomposition of equilateral triangles [16]. Hebert presented a location code for longest-edge bisection hierarchical tetrahedral meshes and a set of rules to compute neighbors efficiently in 3-space [11]. Lee, et al. developed an alternative location code for this same tetrahedral mesh, and presented algorithms for efficient neighbor computation [20]. Both approaches are quite specific to 3-space, and are not readily generalizable to higher dimensions.

We introduce a new location code, which pro-

vides unique encoding of the simplices generated by Maubach’s [6] bisection algorithm. Unlike Hebert’s approach, which only works in 3-dimensional space and relies on enumerations and look-up tables, our approach is fully algorithmic and works in any dimension. All the geometry of the simplices and the operations on the tree can be computed easily based solely on the code of a simplex. Our location code and the definitions of various operations on the simplex tree depend on the particular vertex ordering. We have adopted a different ordering than Maubach’s system, which we feel leads to simpler formulas. Our vertex ordering is a generalization of the vertex ordering used in Hebert’s 3-dimensional system.

The most challenging operation on the tree is neighbor computation. Maubach’s system computes the neighbors of a simplex during construction of the tree recursively [21], and stores pointers to neighbors for each simplex. In contrast, we show how to compute an arbitrary neighbor of any simplex efficiently directly from its code, without storing any neighbor links, and without having to traverse the path to and from the root in order to compute neighbors. This is significant gain both in terms of storage, and computational efficiency, since our approach is local and runs in $O(d)$ time. In fact it runs in $O(1)$ time, if the operations are encoded in lookup tables.

Our interest is in hierarchical regular meshes, however, note that there is also significant interest in compact representation of irregular simplicial meshes [22, 23].

3. PRELIMINARIES

Throughout, we consider real d -dimensional space, \mathfrak{R}^d . We assume that the domain of interest has been scaled to lie within a unit *reference hypercube* of side length 2, centered at the origin, that is $[-1, 1]^d$. We shall denote points in \mathfrak{R}^d using lower-case bold letters, and represent them as d -element row vectors, that is, $\mathbf{v} = (v_1, v_2, \dots, v_d) = (v_i)_{i=1}^d$. We let \mathbf{e}_i denote the i th unit vector. A d -simplex is represented as a $(d + 1) \times d$ matrix whose rows are the vertices of the simplex, numbered from 0 to d . Of particular interest is the *base simplex*, denoted S_\emptyset , whose i th vertex is $\sum_{j=1}^i \mathbf{e}_j - \sum_{j=i+1}^d \mathbf{e}_j$.

Permutations and Reflections: Let $\text{Sym}(d)$ denote the *symmetric group* of all $d!$ permutations over $\{1, 2, \dots, d\}$. We denote a permutation $\Pi \in \text{Sym}(d)$ by a tuple of distinct integers $[\pi_1 \ \pi_2 \ \dots \ \pi_d]$, where $\pi_i \in \{1, 2, \dots, d\}$. We can interpret such a permutation as a linear function that maps the unit vector \mathbf{e}_i to the \mathbf{e}_{π_i} , or equivalently as a coordinate permutation given by a $d \times d$ matrix whose i th row is the unit

vector \mathbf{e}_{π_i} . For example, for $\Pi = [2 \ 3 \ 1]$, $S_\emptyset \Pi =$

$$\begin{bmatrix} -1 & -1 & -1 \\ 1 & -1 & -1 \\ 1 & 1 & -1 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} -1 & -1 & -1 \\ -1 & 1 & -1 \\ -1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}.$$

It is well known that the collection of simplices $\{S_\emptyset \Psi : \Psi \in \text{Sym}(d)\}$ fully subdivides the reference hypercube, and further that this subdivision is compatible [24]. These $d!$ simplices form the starting point of our hierarchical simplicial mesh. The *composition* of two permutations $\Pi \circ \Psi$, defined as $S(\Pi \circ \Psi) = (S\Psi)\Pi$ is given by the matrix product $\Psi\Pi$. Note that the notation $[2 \ 3 \ 1]$ is not a vector in \mathfrak{R}^d , but merely a convenient shorthand for a permutation matrix. Throughout, vectors will be denoted with parentheses, and square brackets will be used for objects that are to be interpreted as linear transformations, or equivalently a shorthand for a matrix. Another useful class of transformations are coordinate reflections, which can be expressed as a d -tuple $R = [r_1 \ r_2 \ \dots \ r_d]$ where $r_i \in \{\pm 1\}$, and is interpreted as a linear transformation represented by the diagonal matrix $\text{diag}(r_1, r_2, \dots, r_d)$.

It will simplify notation to combine the composition of a permutation and a reflection using a unified notation. We define a *signed permutation* to be a d -tuple of integers $[r_i \pi_i]_{i=1}^d$, where $[\pi_i]_{i=1}^d$ is a permutation and $[r_i]_{i=1}^d$ is a reflection. This is interpreted as a linear transformation that maps the i th unit vector to $r_i \mathbf{e}_{\pi_i}$. For example, in \mathfrak{R}^3 , the composition of the reflection $R = [-1 \ -1 \ +1]$ and the permutation $\Pi = [2 \ 3 \ 1]$ is expressed as the signed permutation $[-2 \ -3 \ +1]$, which is just a shorthand for the matrix product $R\Pi$, that is $R\Pi =$

$$\begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & +1 \end{bmatrix} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & -1 \\ +1 & 0 & 0 \end{bmatrix}.$$

The Simplex Decomposition Tree: Recall that the initial simplicial complex is formed from the $d!$ permutations of the base simplex, that is, $S_\emptyset \Psi$ for $\Psi \in \text{Sym}(d)$. Simplices are then refined by a process of repeated subdivision, called *bisection* [6]. (Details will be given below.) The resulting *child* simplices are labeled 0 and 1. By applying the process repeatedly, each simplex in this hierarchy is uniquely identified by its *path*, which is a string over $\{0, 1\}$. The resulting collection of trees is called the *simplex decomposition tree*. It consists of $d!$ separate binary trees, which conceptually are joined under a common super-root. Each simplex of this tree is uniquely identified by a *permutation-path pair* as $S_{\Psi,p}$, where Ψ is the initial permutation of the base simplex, and $p \in \{0, 1\}^*$ is the path string. When starting with the base simplex (Ψ

is the identity permutation) we may omit explicit reference to Ψ . By symmetry, it suffices to describe the bisection process on just the base simplex S_0 . The ordering of the rows, that is, the numbering of vertices, is significant.

Maubach [6] showed that with every d consecutive bisections, the resulting simplices are similar copies of their d -fold grandparent, subject to a uniform scaling by $1/2$. Thus, the pattern of decomposition repeats every d levels in the decomposition. Define the *level*, ℓ , of a simplex S_p to be the path length modulo the dimension, that is, $\ell = (|p| \bmod d)$, where $|p|$ denotes the length of p . The 0-child S_{p0} and 1-child S_{p1} of a simplex are computed as follows:

$$S_p = \begin{bmatrix} \mathbf{v}_0 \\ \dots \\ \mathbf{v}_{\ell-1} \\ \mathbf{v}_\ell \\ \mathbf{v}_{\ell+1} \\ \dots \\ \mathbf{v}_d \end{bmatrix} \quad S_{p0} = \begin{bmatrix} \mathbf{v}_0 \\ \dots \\ \frac{\mathbf{v}_{\ell-1} + \mathbf{v}_d}{2} \\ \mathbf{v}_{\ell+1} \\ \dots \\ \mathbf{v}_d \end{bmatrix} \quad S_{p1} = \begin{bmatrix} \mathbf{v}_0 \\ \dots \\ \frac{\mathbf{v}_{\ell-1}}{2} \\ \mathbf{v}_\ell \\ \dots \\ \mathbf{v}_{d-1} \end{bmatrix}$$

A portion of the tree is illustrated in Fig. 1.

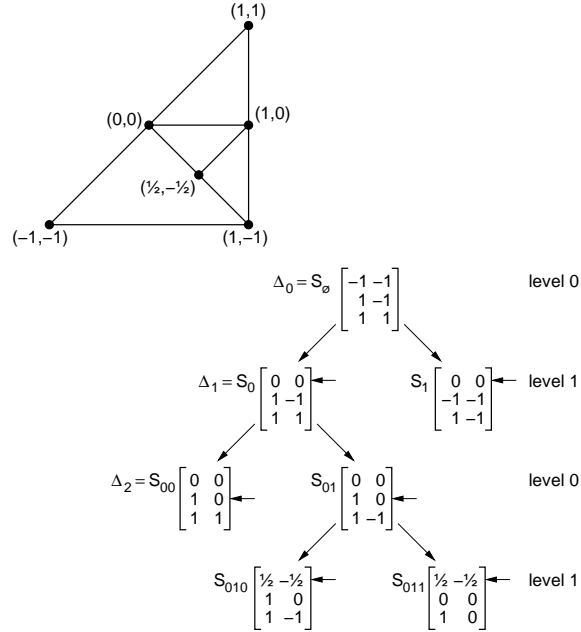


Figure 1: The simplex decomposition tree. The corresponding bisected simplex is shown on top left. The newly created vertex is indicated by an arrow in each case. The reference simplices Δ_i are indicated as well.

Reference Simplices and the Reference Tree: Since with every d consecutive bisections, the simplices are similar to, but half the size, of their d -fold grandparent, we can partition the nodes of the decomposition tree into a collection of isomorphic, disjoint sub-

trees of height d . The roots of these subtrees are the nodes whose depths are multiples of d (where the root starts at depth 0). It suffices to analyze the structure of just one of these trees, in particular, the subtree of height d starting at the root. We call this the *reference tree*. Since the two children of any simplex are congruent, it follows that all the simplices at any given depth of the decomposition tree are congruent to each other. Thus, all the similarity classes are represented by d canonical simplices, called the *reference simplices*. These are defined to be $S_{(0^k)}$, for $0 \leq k < d$, and denoted by Δ_k . (See Fig. 1.) Although it is not a reference simplex, we also define $\Delta_d = S_{(0^d)}$, since it is useful in our proofs.

4. THE LPT code

So far we have defined an infinite decomposition tree and a procedure for generating the simplices of this tree top down. In order to provide pointerless implementation of the hierarchical mesh, we define a *location code*, which uniquely encodes each simplex of the hierarchy. The most direct location code is combination consisting of the initial permutation Ψ followed by the binary encoding of the tree path p . Unfortunately, it is not easy to compute basic properties of the simplex such as neighbors from this code. Instead we modify an approach presented by Hebert [11] for the 3-dimensional case, by defining a location code that more directly encodes the geometric relationship between each simplex and the reference simplex at the same level. We call this the *LPT code*, since it encodes for each simplex its *Level*, its signed *Permutation*, and its *Translation* relative to some reference simplex.

Given any simplex $S_{\Psi,p}$ in the hierarchy, the *LPT code* is a 3-tuple (ℓ, Π, Φ) , where $\ell = |p| \bmod d$ is the simplex's level, Π is a signed permutation relating $S_{\Psi,p}$ to its reference simplex, and Φ is a list of vectors, called the orthant list, which is used to derive the translation relative to the reference simplex. The permutation part $\Pi = \Pi_{\Psi,p}$ and orthant list $\Phi = \Phi_{\Psi,p}$ are defined below as functions of Ψ and p . Correctness will be established in Theorem 4.1 below.

Permutation Part: The signed permutation $\Pi_{\Psi,p}$ is defined recursively as follows for a base permutation Ψ and binary path p :

$$\Pi_{\Psi,\emptyset} = \Psi \quad \Pi_{\Psi,p0} = \Pi_{\Psi,p} \quad \Pi_{\Psi,p1} = \Pi_{\Psi,p} \circ \Sigma_\ell, \quad (1)$$

where Σ_ℓ is the permutation that cyclically shifts the last $d - \ell$ elements to the right and negates the element that is wrapped around. That is, $\Sigma_\ell = [1 \ 2 \ \dots \ \ell \ (-d) \ (\ell + 1) \ (\ell + 2) \ \dots \ (d - 1)]$. A portion of the simplex decomposition tree, and the associated permutation values are shown in Fig. 2. For example, observe that S_1 is related to Δ_1 by the signed permu-

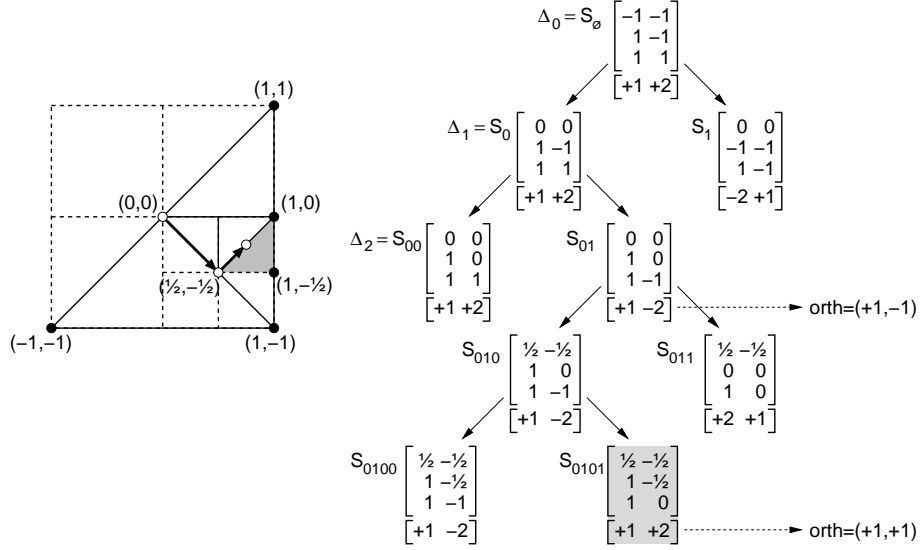
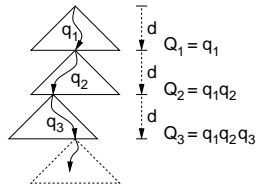


Figure 2: The signed permutations $\Pi_{\Psi,p}$ associated with each simplex are shown below each simplex matrix, and the entries of the orthant list are shown for the shaded simplex S_{0101} . The LPT code for this simplex is $(0, [+1 +2], \langle (+1, -1), (+1, +1) \rangle)$.

tation $[-2 +1]$, which negates the first column of Δ_1 and then swaps the two columns.

Orthant List: Recall that with every d levels of descent in the decomposition tree, the resulting simplices decrease in size by a factor of $1/2$. The bounding hypercube of the resulting descendent is one of the 2^d hypercubes that would result from a quadtree-like decomposition (indicated by broken lines on the left side of Fig. 2). Depending on the level within the tree, the translation of the descendent hypercube relative to its ancestor will be some power of $(1/2)$ times a d -vector over $\{\pm 1\}$. Such a vector defines the *orthant* containing the descendent hypercube relative to the central vertex of its ancestor. Consider the shaded simplex in Fig. 2. Its translation relative to the base simplex is $\frac{1}{2}(+1, -1) + \frac{1}{4}(+1, +1)$, indicated by the arrowed lines. The orthant list encodes these two vectors.



To define the orthant list, we first remove the last ℓ symbols of p , leaving a multiple of d symbols (possibly empty). We then partition the remaining symbols into $L = \lfloor |p|/d \rfloor$ substrings, $q_1 q_2 \dots q_L$, where $|q_i| = d$. Since the reference tree structure repeats

every d levels, each q_i can be viewed as a complete path in one of these subtrees of height d . Let Q_i denote the concatenation of the first i substrings. For $1 \leq i \leq L$, define $\Gamma_{\Psi,p}[i]$ to be the signed permutation for path Q_i , that is Π_{Ψ,Q_i} . Given a signed permuta-

tion Γ , define $orth(\Gamma)$ to be the d -vector that is the image of the row d -vector $\mathbf{1}_d$ under Γ . For example, if $\Pi = [-2 -3 +1]$ then $orth(\Pi) = (+1, -1, -1)$. Define the *orthant list* for the pair (Ψ, p) to be the sequence of L vectors whose i th element is $orth(\Gamma_{\Psi,p}[i])$, that is

$$\Phi_{\Psi,p} = \langle orth(\Gamma_{\Psi,p}[1]), orth(\Gamma_{\Psi,p}[2]), \dots, orth(\Gamma_{\Psi,p}[L]) \rangle \quad (2)$$

The orthant list can be computed incrementally along with the permutation part of the code as follows. Given the LPT code (ℓ, Π, Φ) for a simplex $S_{\Psi,p}$, first observe that the orthant list only changes for the children if the current level is $d - 1$. If so, we compute the child's permutation Π' from Eq. (1) and append $orth(\Pi')$ to the current list.

```

LPTcode(p, (l, Pi, Phi))
  if (p = empty) return (l, Pi, Phi)
  Express p as xq, for x in {0, 1}
  l ← (l + 1) mod d
  if (x = 1) Pi ← Pi ∘ Σ_l
  if (l = 0) Phi ← Phi + orth(Pi)
  return LPTcode(q, (l, Pi, Phi))

```

The computation of the LPT code is summarized in the procedure *LPTcode*. The code for the simplex $S_{\Psi,p}$ is computed by the call *LPTcode*($p, (0, \Psi, \emptyset)$). We may now state the main result of this section, called the *LPT Theorem*, which establishes the geometric meaning of our LPT code by relating each simplex of the decomposition tree to its associated reference simplex. Hebert [11] proved the analogous result for his 3-dimensional bisection system. Let $\mathbf{1}_{d+1}^T$ denote a $(d + 1)$ -column vector of 1's. The following theo-

rem makes use of the observation that, for any d -row vector \mathbf{v} , the matrix product $\mathbf{1}_{d+1}^T \cdot \mathbf{v}$ is a $(d+1) \times d$ vector whose rows are all equal to \mathbf{v} , thus adding this to a simplex matrix is equivalent to a translation by \mathbf{v} . The proof is given in the appendix.

Theorem 4.1 (LPT Theorem) *Let $S_{\Psi,p}$ be the simplex of the decomposition tree associated with some initial permutation Ψ and binary path p . Let (ℓ, Π, Φ) be the LPT code for this simplex, defined above. Then $S_{\Psi,p}$ is related to Δ_ℓ , the reference simplex at this level, by the following similarity transformation:*

$$S_{\Psi,p} = \frac{1}{2^L} \Delta_\ell \Pi + \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi[i].$$

where $L = \lfloor |p|/d \rfloor$.

We can now describe a pointerless implementation of a simplex decomposition tree. For each simplex $S_{\Psi,p}$ in the tree, we create a node that is indexed by an appropriate encoding of the associated LPT code. Theorem 4.1 implies that the geometry of this simplex is determined entirely from the LPT code, and, if desired, it can be computed from the code in time proportional to the code length. These objects are then stored in any index structure that supports rapid look-ups, for example, a hash table.

Implementation Issues: There are a number of practical observations that can be made on how to encode LPT codes efficiently in low dimensional spaces. Let D denote the maximum depth of any node in the tree. Each of the $d!$ permutations of $\text{Sym}(d)$ can be encoded as an integer with $\log_2 d!$ bits [25]. A d -element reflection vector over $\{\pm 1\}$ can be represented as a d -element bit string (e.g., by the mapping $+1 \rightarrow 0$ and $-1 \rightarrow 1$). Thus, a signed permutation Π then can be encoded by a pair of integers. A convenient way to encode the vectors of the orthant list is map them to bit strings and to store them as d separate lists, one for each coordinate. (The advantage of this representation will be discussed in Section 5.3.) The final code consists of the level ℓ , expressed with $\lceil \log_2 d \rceil$ bits, the permutation and reflection, represented using $\lceil \log_2(d!) \rceil + d$ bits, and the orthant list, represented using $d \cdot \text{length}(\Phi)$ bits, which is at most $d \lfloor D/d \rfloor \leq D$. The total number of bits needed to represent the code for a simplex at depth D is $D + \log_2(d!) + O(d)$. This is close to optimal in the worst case, since there are $2^D d!$ simplices at depth D in a full tree. If we assume that the machine's word size is $\Omega((D/d) + \log_2 d!)$, then the permutation part of the code can be stored in a constant number of machine words and the orthant lists can be stored in $O(d)$ machine words.

Also, note that for small d , the multiplication tables for the various signed permutations (such as Σ_ℓ of Eq. (1) and the neighbor permutations of Section 5.3 below) can be precomputed and stored in tables. This

allows very fast evaluation of permutation operations by simple table look-up.

5. DECOMPOSITION TREE OPERATIONS

In this section we present methods for performing useful tree access operations based on manipulations of LPT codes.

5.1 Tree Traversal

Consider a simplex $S_{\Psi,p}$ of the tree whose LPT code is (ℓ, Π, Φ) . Let us consider how to compute the children and parent of this simplex in the tree. The LPT codes of the children of this simplex can be computed in $O(d)$ time by applying the recursive rules used to define the LPT code, given in Section 4. We can compute the parent from the LPT code by inverting this process, but to do so we need to know whether the simplex is a 0-child, a 1-child, or the root. A root simplex is distinguished by having an empty orthant list and level $\ell = 0$. Otherwise, we use the following lemma.

Lemma 5.1 *Consider a nonroot simplex S of the decomposition tree with LPT code (ℓ, Π, Φ) , and let S' be its nearest proper ancestor at level 0. Let $\Pi = [\pi_i]_{i=1}^d$ be the signed permutation of S , let $\mathbf{o} = (o_i)_{i=1}^d$ be the last entry of the orthant list of S' , and let $\ell^* = 1 + ((\ell - 1) \bmod d)$. Then S is a 0-child if and only if $\text{sign}(\pi_{\ell^*}) = \text{sign}(o_{|\pi_{\ell^*}|})$.*

This lemma can be applied to determine the LPT code for the parent of a nonroot simplex S . Given S 's LPT code, (ℓ, Π, Φ) , we distinguish two cases, depending on its level. If ℓ is nonzero, then its parent's level is $\ell' = \ell - 1$ and otherwise its parent's level is $\ell' = d - 1$. If ℓ is nonzero, then the orthant vector \mathbf{o} of the lemma is the last entry of Φ . We apply this lemma to determine whether S is a 0- or 1-child. From Eq. (1) and Theorem 4.1 we know that, if it is a 0-child, it has the same permutation code as its parent, and otherwise its parent's permutation code is $\Pi \circ \Sigma_{\ell'}^{-1}$. Its parent has the same orthant list. On the other hand, if $\ell = 0$ then \mathbf{o} is the second to last entry of Φ . Again we apply the lemma to determine whether S is a 0- or 1-child, and derive its parent's permutation code. The last entry of S 's orthant list is removed to form the orthant list of its parent. This can be computed in $O(d)$ time.

5.2 Point Location and Interpolation

In this section we consider how to compute the LPT code of the leaf simplex of the decomposition tree that contains a given query point $\mathbf{q} = (q_i)_{i=1}^d$. We assume that \mathbf{q} lies in the base hypercube, that is, $-1 \leq q_i \leq 1$. If \mathbf{q} lies on a face between two simplices, we will choose one arbitrarily.

We begin by locating the root simplex, $S_{\Psi,0}$ that contains \mathbf{q} . It is easy to see that a point \mathbf{q} in the base

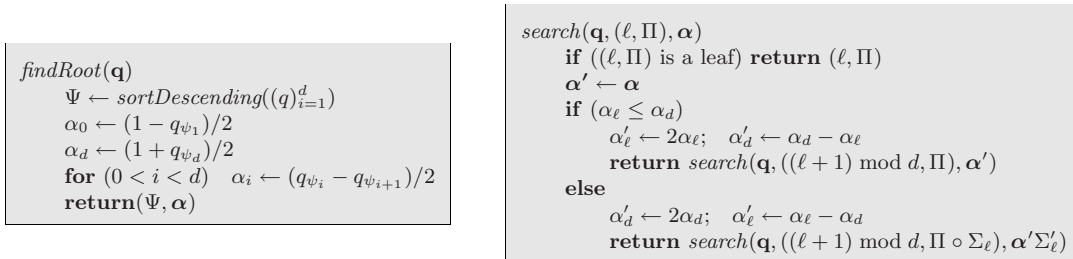


Figure 3: The procedures *findRoot* and *search*, which are used to locate a query point \mathbf{q} in the hierarchy. The permutation Σ'_ℓ is defined in Lemma 5.2 and the permutation Σ_ℓ was given in Section 4, Eq. 1.

hypercube lies in the base reference simplex, Δ_0 , if and only if its coordinate vector is sorted in decreasing order. It follows that determining the permutation Ψ of the root simplex reduces to sorting the coordinates of \mathbf{q} in decreasing order and setting Ψ to the permutation that produces this sorted order. Let us assume that the function *sortDescending* computes this permutation. Letting \mathbf{v}_i denote the i th vertex of the root simplex $S_{\Psi, \emptyset}$ that contains \mathbf{q} , the *barycentric coordinates* of \mathbf{q} with respect to $S_{\Psi, \emptyset}$ is the unique $d + 1$ vector $\alpha = (\alpha_i)_{i=0}^d$, $0 \leq \alpha_i \leq 1$, such that $\sum_i \alpha_i = 1$ and $\mathbf{q} = \sum_i \alpha_i \mathbf{v}_i$. Because of the special structure of Δ_0 , it is easy to verify that the procedure *findRoot* shown in Fig. 3 computes these coordinates.

After this initialization, we recursively descend the hierarchy until finding a leaf simplex. We use the barycentric coordinates of \mathbf{q} relative to the current simplex to determine in which child it resides. Then we generate the barycentric coordinates of \mathbf{q} with respect to this child. This is done with the aid of the following lemma. The descent algorithm is given in Fig. 3. To simplify the presentation, we have omitted the orthant list processing, which is the same as in the code block just prior to Theorem 4.1.

Lemma 5.2 *Consider a nonleaf simplex $S_{\Psi, p}$ of the hierarchy at level ℓ with the associated permutation code $\Pi_{\Psi, p} = [\pi_i]_{i=1}^d$. Suppose that \mathbf{q} lies within this simplex with the barycentric coordinates $\alpha = (\alpha_i)_{i=0}^d$.*

- *If $\alpha_\ell \leq \alpha_d$, then \mathbf{q} lies in the 0-child. Let α' be the $(d + 1)$ -vector that is identical α except that $\alpha'_\ell = 2\alpha_\ell$ and $\alpha'_d = \alpha_d - \alpha_\ell$. Then the barycentric coordinate vector of \mathbf{q} relative to this child is α' .*
- *Otherwise, \mathbf{q} lies in the 1-child. Let Σ'_ℓ be a $(d + 1)$ -permutation that shifts the last $d + 1 - \ell$ coordinates circularly one position to the right. Let α' be the $(d + 1)$ -vector that is identical to α except that $\alpha'_d = 2\alpha_d$ and $\alpha'_\ell = \alpha_\ell - \alpha_d$. Then the barycentric coordinate vector of \mathbf{q} relative to this child is $\alpha' \Sigma'_\ell$.*

Given the query point \mathbf{q} , the point location procedure first calls *findRoot* to find the appropriate root simplex

Ψ of the decomposition tree and the barycentric coordinates α . Then it invokes the recursive procedure *search*($0, \Psi, \alpha$) to locate \mathbf{q} within the appropriate root simplex. Once the point has been located, we can answer the interpolation query for this point. We access the stored vector field values at each of the simplex vertices, and then weight these values according to the barycentric coordinates of \mathbf{q} . The result is a piecewise linear, continuous interpolant.

This simple sequential search makes as many memory accesses as the depth of the final leaf simplex that contains \mathbf{q} . A more efficient procedure in terms of memory accesses would be to employ a doubling binary search, which computes (using only local memory) the LPT codes for the simplices at depths 0, 1, 2, 4, 8, and so on, until first finding a depth whose simplex does not exist in the hierarchy. We then use standard binary search to locate the exact depth of the leaf simplex that contains \mathbf{q} . Although the computation of the LPT codes is performed sequentially in time linear in the depth of the final simplex, the number accesses to the simplex decomposition tree is only logarithmic in the final depth. Thus, the running time is $O(dD)$, where D is the maximum depth of the tree, and $O(\log D)$ global memory accesses are made.

5.3 Neighbors in the Simplicial Complex

As we mentioned earlier, when a simplex of the decomposition tree is bisected, it is necessary to bisect some of its neighbors in order to guarantee that the final subdivision is compatible. Thus, it is necessary to compute neighbors within the complex. Two simplices are *neighbors* if they share a common $(d - 1)$ -dimensional face. In addition to this major need, computing facet neighbors of a simplex efficiently is of great interest for many applications that require moving along adjacent simplices, such as direct volume rendering and isosurface extraction techniques.

Consider a simplex S in the complex defined by the decomposition tree. For $0 \leq i \leq d$, let \mathbf{v}_i denote its i th vertex. Exactly one $(d - 1)$ -face of S does not

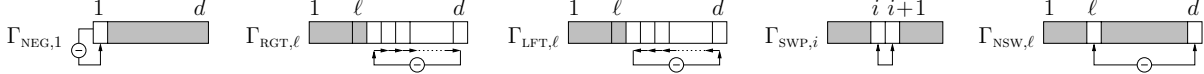


Figure 4: Permutation types. (The circle with a minus sign indicates that the element is negated.)

$$\begin{array}{llll}
\text{if } (S \text{ is a 0-child}) : & N^{(0)}(S) : & & \Pi^{(0)} = \Pi \circ \Gamma_{\text{NEG},1} \\
& N^{(i)}(S) : \quad (0 < i < d) & & \Pi^{(i)} = \Pi \circ \Gamma_{\text{SWP},i} \\
& N^{(d)}(S) : & & \Pi^{(d)} = \Pi \circ \Gamma_{\text{RGT},\ell^-} \\
\text{if } (S \text{ is a 1-child}) : & N^{(0)}(S) : & & \Pi^{(0)} = \Pi \circ \Gamma_{\text{NEG},1} \\
& N^{(\ell^*)}(S) : & & \Pi^{(\ell^*)} = \Pi \circ \Gamma_{\text{LFT},\ell^-} \\
& N^{(i)}(S) : \quad (0 < i < d, i \neq \ell^*) & & \Pi^{(i)} = \Pi \circ \Gamma_{\text{SWP},i} \\
& N^{(d)}(S) : \quad (d \neq \ell^*) & & \Pi^{(d)} = \Pi \circ \Gamma_{\text{NSW},\ell}
\end{array}$$

Figure 5: Permutation codes for neighboring simplices.

contain \mathbf{v}_i . If this face is not on the boundary of the base hypercube, its neighbor exists in the complex. If so, we define $N^{(i)}(S)$ to be the same depth neighboring simplex to S lying on the opposite side of this face. Let (ℓ, Π, Φ) denote the LPT code for S and let $(\ell^{(i)}, \Pi^{(i)}, \Phi^{(i)})$ denote the LPT code for $N^{(i)}(S)$. We present rules here for computing LPT codes of these neighbors. The proof of their correctness is based on a straightforward but lengthy induction argument. A sketch of the proof is presented in the appendix. The rules compute the LPT code for the neighbor simplex at the same depth as S , and hence $\ell^{(i)} = \ell$. Of course, this simplex need not be in the decomposition tree because its parent may not yet have been bisected. In fact, in a compatible subdivision, a $(d-1)$ -face neighbor of S could also appear at one level higher or one level lower than S . We will also show how to compute the LPT codes of those neighbors.

Neighbor Permutation Code: Each neighbor's permutation code is determined by applying one of a set of special signed permutations to Π . The permutation depends on whether S is a 0-child or a 1-child, which can be determined using the test given in Section 5.1. These permutations are illustrated in Fig. 4, and include the following: $\Gamma_{\text{NEG},1}$, negates the first element, $\Gamma_{\text{RGT},\ell}$ (resp., $\Gamma_{\text{LFT},\ell}$), shifts the last $d-\ell$ elements cyclically one position to the right (resp., left) and negates the element that was wrapped around, $\Gamma_{\text{SWP},i}$, swaps elements i and $i+1$, and $\Gamma_{\text{NSW},\ell}$, swaps and negates elements ℓ and d . The neighbor rules are given in Fig. 5. A number of the rules involve the parent's level, and so to condense notation, we define $\ell^- = (\ell-1) \bmod d$ and $\ell^* = \ell^- + 1$. Observe that $\ell^- = \ell-1$ and $\ell^* = \ell$, except when $\ell=0$, in which case they are larger by d . These can be computed in $O(d)$ time, and in fact in $O(1)$ time if permutations are encoded in look-up tables.

Neighbor Orthant: If $\ell \neq 0$ or $1 \leq i < d$, $N^{(i)}(S)$, is in the same final orthant as S , and so $\Phi^{(i)} = \Phi$. If

$\ell=0$, $N^{(d)}$ is in a different orthant than S , but, $N^{(d)}$ is the sibling of S in this case. Thus, Φ and $\Phi^{(d)}$ differ only in their last element, which is $\text{orth}(\Pi^{(d)})$. Thus the orthant list can be updated in either $O(1)$ or $O(d)$ time in these cases.

The only remaining case is $\Phi^{(0)}$. This case is the most complex because the final enclosing quadtree box of $N^{(0)}(S)$ is disjoint from S 's final quadtree box. Further, it may be arbitrarily far away, in the sense that the least common ancestor of the two nodes may be the root of the tree. In this case, to compute $\Phi^{(0)}$, we can apply known methods for computing neighbors in linear quadtrees [15], which provide the answer to the following problem: Given the path from the root to an orthant A (in our representation this path corresponds to the orthant list), and a direction defined as a 2-tuple (D, X_i) where X_i is the i^{th} coordinate axis, and $D \in \{-, +\}$ represents the direction of X_i , find the path to the neighbor orthant of equal size located in the given direction with respect to A . Thus, in order to compute $\Phi^{(0)}$, which represents the path from the root to the final orthant of $N^{(0)}(S)$, all we need is to determine which direction $N^{(0)}(S)$ is located with respect to S . Consider the Π and $\Pi^{(0)}$ corresponding to S and $N^{(0)}(S)$ respectively. By the given neighbor rules, these two permutation-reflection codes differ only in the sign of their first element. This is the sign corresponding to the $X_{|\pi_1|}$ axis, given that $\Pi = [\pi_i]_1^d$ is the code for S . The sign of π_1 determines in which direction of $X_{|\pi_1|}$ axis S resides in its final orthant. And so, the neighbor $N^{(0)}(S)$ is also in that direction. Thus, the axis component of the direction is $X_{|\pi_1|}$, and the sign component of the direction is $\text{sign}(\pi_1)$.

Implementation Issues: We omit the details, due to space limitations, except to note that the operation can be implemented very rapidly through a simple trick with bit manipulations. The neighbor computation [15] essentially involves an operation, which

is applied to a bit string that consists of the i th coordinate of each entry of the orthant list. Recall that the orthant list is stored as d separate bit strings, one per coordinate, and packed into machine words as binary numbers. The key operation needed for the neighbor computation involves complementing a maximal trailing sequence of matching bits. For example, given a bit string of the form $w10^k$, for $w \in \{0, 1\}^*$, the desired result is $w01^k$ (and vice versa). By packing these bits into a single word, we can compute this function with a single arithmetic operation by subtracting (or adding) 1 from the resulting binary number. Assuming that the machine's word size is $\Omega(D/d)$, where D is the maximum depth of any simplex, the orthant list for the neighbor can be computed in $O(1)$ time.

Neighbors at different depths: Above rules provide the LPT code for the same depth neighbors. However, in a compatible subdivision, a neighbor could possibly appear one level closer or one level further from the root, that is, some neighbors of a simplex S_p at depth $|p|$, could appear at depths $|p| - 1$ or $|p| + 1$. We can categorize the neighbors of a simplex into two groups: neighbors that share the edge-to-be-bisected, and neighbors that do not. Maubach already proved that a neighbor sharing the edge-to-be-bisected is either at depth $|p|$ or at depth $|p| - 1$, and that a neighbor at depth $|p| - 1$ is the parent of the same depth neighbor which did not come into existence yet. Hence, to compute the LPT code for a neighbor at depth $|p| - 1$, we first compute the LPT code for the same depth neighbor by the above rules, and if the same depth neighbor does not exist in the tree, we compute its parent's LPT code as described in Section 5.

In addition, any $(d-1)$ -face neighbor of S_p which does not share the edge-to-be-bisected could possibly be at depth $|p| + 1$. Specifically, same depth neighbors $N^{(\ell)}(S_p)$ and $N^{(d)}(S_p)$, might have been bisected without triggering bisection of S_p , and so, one of their children will now share a face with S_p . Moreover, the child of $N^{(\ell)}(S_p)$ or $N^{(d)}(S_p)$ that shares a face with S_p , is the same depth neighbor of one of the children of S_p . So, we can compute a neighbor at depth $|p| + 1$ by computing the appropriate same depth neighbor of one of the children of S_p . Formally,

if($N^{(\ell)}(S_p)$ is a bisected simplex)
 $N^{(\ell)}(S_{p0})$ is the ℓ^{th} neighbor of S_p .
 if($N^{(d)}(S_p)$ is a bisected simplex)
 $N^{(\ell)}(S_{p1})$ is the d^{th} neighbor of S_p .

It can be easily proven that these neighbors cannot exist at depths higher than $|p| + 1$.

6. CONCLUSION

We have presented a representation of hierarchical regular simplicial meshes based on Maubach's [6] bisection

algorithm. Unlike Maubach's approach, which requires the use of recursion or an explicit tree structure, our representation is pointerless, that is, the simplices of the mesh are uniquely identified through a location code, called the LPT code. We have shown how to use this code to traverse the hierarchy, compute neighbors, and to answer point location and interpolation queries.

References

- [1] Bank R., Sherman A., Weiser A. "Refinement algorithms and data structures for regular local mesh refinement." *Sci. Comp.*, pp. 3–17, 1983
- [2] Mitchell W.F. "Optimal Multilevel Iterative Methods for Adaptive Grids." *SIAM J. Sci. Stat. Comp.*, vol. 13, 146–167, 1992
- [3] Arnold D., Mukherjee A., L.Pouly. "Locally adapted tetrahedral meshes using bisection." *SIAM J. Sci. Comp.*, vol. 22, no. 2, 431–448, 2001
- [4] Bey J. "Tetrahedral Grid Refinement." *Computing*, vol. 55, 355–378, 1995
- [5] Liu A., Joe B. "Quality local refinement of tetrahedral meshes based on 8-subtetrahedron subdivision." *Math. of Comput.*, vol. 65, no. 215, 1183–1200, 1996
- [6] Maubach J.M. "Local bisection refinement for N -simplicial grids generated by reflection." *SIAM J. Sci. Stat. Comp.*, vol. 16, 210–227, 1995
- [7] Rivara M. "Local modification of meshes for adaptive and/or multigrid finite-element methods." *J. Comp. Appl. Math.*, vol. 36, 79–89, 1991
- [8] Munkres J.R. *Topology: A first course*. Prentice Hall, Englewood Cliffs, NJ, 1975
- [9] Zhou Y., Chen B., Kaufman A. "Multiresolution tetrahedral framework for visualizing regular volume data." *Proc. IEEE Visualization'97*, pp. 135–142. 1997
- [10] Gerstner T., Rumpf M. "Multiresolutional parallel isosurface extraction based on tetrahedral bisection." *Proc. Symp. Volume Visualization*. 1999
- [11] Hebert D.J. "Symbolic local refinement of tetrahedral grids." *J. of Symbolic Comp.*, vol. 17, 457–472, 1994
- [12] Glassner A.S. *An Introduction to Ray Tracing*. Academic Press, San Diego, 1989
- [13] Atalay F.B., Mount D.M. "Ray interpolants for fast ray-tracing reflections and refractions." *J. of WSCG*, vol. 10, no. 3, 1–8, 2002. Proc. Int. Conf. in Central Europe on Comp. Graph., Visual. and Comp. Vision
- [14] Atalay F.B., Mount D.M. "Interpolation over light fields with applications in computer graphics." *Proc. 5th Workshop on Algorithm Eng. and Exp. (ALENEX 2003)*, pp. 56–68. SIAM, 2003
- [15] Samet H. *Applications of Spatial Data Structures: Computer Graphics, Image Processing, and GIS*. Addison-Wesley, 1990

- [16] Lee M., Samet H. “Navigating through triangle meshes implemented as linear quadtrees.” *ACM Trans. on Comput. Graph.*, vol. 19, 79–121, 2000
- [17] Chilimbi T., Hill M.D., Larus J.R. “Cache-Conscious Structure Layout.” *Programming Languages Design and Implementation*. 1999
- [18] Aggarwal A., Vitter J.S. “The input/output complexity of sorting and related problems.” *Commun. ACM*, vol. 31, 1116–1127, 1988
- [19] Gargantini I. “An Effective Way to Represent Quad-trees.” *Commun. ACM*, vol. 25, no. 12, 905–910, 1982
- [20] Lee M., Floriani L.D., Samet H. “Constant-time neighbor finding in hierarchical tetrahedral meshes.” *Proc. Int. Conf. on Shape Modelling*, pp. 286–295. 2001
- [21] Maubach J.M. “The efficient location of neighbors for locally refined n -simplicial grids.” *5th Int. Meshing Roundtable*. 1996
- [22] Blandford D.K., Blesch G.E., Cardoze D.E., Kadow C. “Compact Representations of Simplicial Meshes in two and three dimensions.” *Proc. 12th International Meshing Roundtable*. 2003
- [23] Szymczaka A., Rossignac J. “Grow & Fold: compressing the connectivity of tetrahedral meshes.” *Comp. Aided Design*, vol. 32, 527–537, 2000
- [24] Allgower E., Georg K. “Generation of Triangulations by Reflection.” *Utilitas Mathematica*, vol. 16, 123–129, 1979
- [25] Knuth D.E. *Sorting and Searching*, vol. 3 of *The Art of Computer Programming*. 1973
- [26] Atalay F.B., Mount D.M. “Pointerless implementation of hierarchical simplicial meshes and efficient neighbor finding in arbitrary dimensions.” Tech. Rep. CS-TR-4586/UMIACS-TR-2004-29, University of Maryland, College Park, 2004

APPENDIX

Due to space limitations, we have omitted some of the cases in the following proofs and proofs of the lemmas. For complete proofs, please refer to the full paper [26].

Proof of Theorem 4.1

We define the following functions that act on a signed permutation $\Pi = [\pi_i]_{i=1}^d$. The first, $perm(\Pi)$, extracts the permutation part of Π , the second, $refl(\Pi)$, extracts the (unpermuted) reflection part as a vector in $\{\pm 1\}^d$. More formally,

$$perm(\Pi) = [|\pi_i|]_{i=1}^d \quad refl(\Pi) = (sign(\pi_i))_{i=1}^d$$

Note also that $orth(\Pi) = refl(\Pi)perm(\Pi)$. The following is an easy consequence.

Lemma 6.1 *Let Π be a signed permutation. Then $\Pi = diag(refl(\Pi))perm(\Pi) = perm(\Pi)diag(orth(\Pi))$.*

Inductive Hypothesis: Let $S_{\Psi,p}$ denote any simplex. Let $\ell = |p| \bmod d$ and let Δ_ℓ be the reference simplex for this level. Then $S_{\Psi,p}$ is similar to Δ_ℓ , and in particular

$$S_{\Psi,p} = \frac{1}{2^L} \Delta_\ell \Pi_{\Psi,p} + \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p}[i].$$

where $L = \lfloor |p|/d \rfloor$, and the signed permutation $\Pi_{\Psi,p}$ and orthant list $\Phi_{\Psi,p}[i]$ are defined in Section 4.

Induction Basis: The hypothesis holds for all root simplices, $S_{\Psi,\emptyset}$, where $L = \lfloor |p|/d \rfloor = 0$, and $\ell = 0$,

$$\begin{aligned} S_{\Psi,\emptyset} &= \frac{1}{2^0} \Delta_0 \Pi_{\Psi,\emptyset} + \mathbf{1}_{d+1}^T \sum_{i=1}^0 \frac{1}{2^i} \Phi_{\Psi,\emptyset}[i] \\ &= \Delta_0 \Pi_{\Psi,\emptyset} \quad (\text{holds by definition}). \end{aligned}$$

Induction Step: Assume that the inductive hypothesis holds for $S_{\Psi,p}$, at level $\ell = |p| \bmod d$. We will show that, it holds for the 0- and 1-children of $S_{\Psi,p}$.

We will make use of the following equalities and lemmas. We can define $S_{p0} = B_{\ell,0} S_p$ and $S_{p1} = B_{\ell,1} S_p$, where $B_{\ell,0}$ and $B_{\ell,1}$ are $(d+1) \times (d+1)$ matrices whose ℓ th row (starting from row 0) has the value $1/2$ in columns ℓ and d (starting from column 0), and all other rows are unit vectors. Thus, $\Delta_{\ell+1} = B_{\ell,0} \Delta_\ell$.

Lemma 6.2 *Given the reference simplex Δ_ℓ , $0 \leq \ell < d$,*

$$B_{\ell,0} \Delta_\ell = B_{\ell,1} \Delta_\ell \Sigma_\ell^{-1},$$

where Σ_ℓ is as defined in Section 4.

Lemma 6.3 *Given the reference simplex Δ_ℓ , $0 \leq \ell < d$, and a signed permutation $\Pi_{\Psi,p}$,*

$$B_{\ell,1} \Delta_\ell \Pi_{\Psi,p} = \Delta_{\ell+1} \Pi_{\Psi,p1}$$

Let $T = \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p}[i]$. Note that, $T = \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p0}[i] = \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p1}[i]$ as well. Also, note that all rows of T are equal to each other. For such a matrix T , the following equalities hold,

$$B_{\ell,0} T = T, \quad B_{\ell,1} T = T.$$

In the induction, there are two cases depending on ℓ :

1. $0 \leq \ell < d-1$

(a) First, consider $S_{\Psi,p0}$. By definition, $\Pi_{\Psi,p0} = \Pi_{\Psi,p}$.

$$\begin{aligned} S_{\Psi,p0} &= B_{\ell,0} S_{\Psi,p} \\ &= B_{\ell,0} \left(\frac{1}{2^L} \Delta_\ell \Pi_{\Psi,p} + T \right) \quad (\text{ind. hyp.}) \\ &= \frac{1}{2^L} \Delta_{\ell+1} \Pi_{\Psi,p} + T \\ &= \frac{1}{2^L} \Delta_{\ell+1} \Pi_{\Psi,p0} + T. \end{aligned}$$

This completes the induction for $S_{\Psi,p0}$, since $\lfloor |p0|/d \rfloor = L$ for $0 \leq \ell < d-1$.

(b) Now, consider S_{Ψ,p_1} .

$$\begin{aligned} S_{\Psi,p_1} &= B_{\ell,1} S_{\Psi,p} \\ &= B_{\ell,1} \left(\frac{1}{2^\ell} \Delta_\ell \Pi_{\Psi,p} + T \right) \quad (\text{ind. hyp.}) \\ &= \frac{1}{2^\ell} \Delta_{\ell+1} \Pi_{\Psi,p_1} + T. \quad (\text{Lemma 6.3}) \end{aligned}$$

This completes the induction for S_{Ψ,p_1} , since $\lfloor |p_1|/d \rfloor = L$ for $0 \leq \ell < d-1$.

2. $\ell = d-1$, the children of $S_{\Psi,p}$ will be at *level 0*.

(a) First, consider S_{Ψ,p_0} . By definition, $\Pi_{\Psi,p_0} = \Pi_{\Psi,p}$.

$$\begin{aligned} S_{\Psi,p_0} &= B_{d-1,0} S_{\Psi,p} \\ &= B_{d-1,0} \left(\frac{1}{2^L} \Delta_{d-1} \Pi_{\Psi,p} + \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p}[i] \right) \\ &= \frac{1}{2^L} \Delta_d \Pi_{\Psi,p} + \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p}[i] \\ &= \frac{1}{2^L} \Delta_d \Pi_{\Psi,p_0} + \mathbf{1}_{d+1}^T \sum_{i=1}^L \frac{1}{2^i} \Phi_{\Psi,p_0}[i] \\ &= \frac{1}{2^L} \Delta_d \Pi_{\Psi,p_0} + \mathbf{1}_{d+1}^T \sum_{i=1}^{L+1} \frac{1}{2^i} \Phi_{\Psi,p_0}[i] \\ &\quad - \frac{1}{2^{L+1}} \mathbf{1}_{d+1}^T \Phi_{\Psi,p_0}[L+1] \end{aligned}$$

Since $\Delta_d = \frac{\Delta_0 + [\mathbf{1}]_{(d+1) \times d}}{2}$ where $[\mathbf{1}]_{(d+1) \times d}$ is a matrix of 1's and $\Phi_{\Psi,p_0}[L+1] = \text{orth}(\Pi_{\Psi,p_0})$, S_{Ψ,p_0} is equal to

$$\begin{aligned} &\frac{1}{2^L} \frac{(\Delta_0 + [\mathbf{1}]_{(d+1) \times d})}{2} \Pi_{\Psi,p_0} + \mathbf{1}_{d+1}^T \sum_{i=1}^{L+1} \frac{1}{2^i} \Phi_{\Psi,p_0}[i] \\ &\quad - \frac{1}{2^{L+1}} \mathbf{1}_{d+1}^T \text{orth}(\Pi_{\Psi,p_0}) \\ &= \frac{1}{2^{L+1}} \Delta_0 \Pi_{\Psi,p_0} + \mathbf{1}_{d+1}^T \sum_{i=1}^{L+1} \frac{1}{2^i} \Phi_{\Psi,p_0}[i] \\ &\quad + \frac{1}{2^{L+1}} ([\mathbf{1}]_{(d+1) \times d} \Pi_{\Psi,p_0} - \mathbf{1}_{d+1}^T \text{orth}(\Pi_{\Psi,p_0})) \end{aligned}$$

By Lemma 6.1,

$$\begin{aligned} \mathbf{1}_{d+1}^T \text{orth}(\Pi_{\Psi,p_0}) &= \mathbf{1}_{d+1}^T \text{refl}(\Pi_{\Psi,p_0}) \text{perm}(\Pi_{\Psi,p_0}) \\ &= [\mathbf{1}]_{(d+1) \times d} \text{diag}(\text{refl}(\Pi_{\Psi,p_0})) \text{perm}(\Pi_{\Psi,p_0}) \\ &= [\mathbf{1}]_{(d+1) \times d} \Pi_{\Psi,p_0}. \end{aligned}$$

We see that, the third term above is 0, yielding

$$S_{\Psi,p_0} = \frac{1}{2^{L+1}} \Delta_0 \Pi_{\Psi,p_0} + \mathbf{1}_{d+1}^T \sum_{i=1}^{L+1} \frac{1}{2^i} \Phi_{\Psi,p_0}[i]$$

This completes the induction for S_{Ψ,p_0} , since $\lfloor |p_0|/d \rfloor = L+1$ for $\ell = d-1$.

(b) Case of S_{Ψ,p_1} can be similarly proven applying the same derivations as in case 1(b) and 2(a).

Correctness of the Neighbor Rules

The following notation will be used throughout the proof. S denotes any simplex. $S^{(i)} = N^{(i)}(S)$, i.e. the i^{th} neighbor of S . Π and $\Pi^{(i)}$ denote the signed permutation code associated with S and $S^{(i)}$ respectively. S_0 and S_1 denote the 0- and 1-children of S , respectively. $S_0^{(i)}$ and $S_1^{(i)}$ denote the 0- and 1-children of $S^{(i)}$, respectively. Π_0 and Π_1 denote the signed permutation code associated with S_0 and S_1 , respectively. $\Pi_0^{(i)}$ and $\Pi_1^{(i)}$ denote the signed permutation code associated

with $S_0^{(i)}$ and $S_1^{(i)}$, respectively. $(S_0)^{(i)}$ denote the i^{th} neighbor of S_0 . $(\Pi_0)^{(i)}$ denotes the code for $(S_0)^{(i)}$. $(S_1)^{(i)}$ denote the i^{th} neighbor of S_1 . $(\Pi_1)^{(i)}$ denotes the code for $(S_1)^{(i)}$. \mathbf{m} and \mathbf{m}' are used to denote the new vertex generated by bisection. \mathbf{u} is used for the vertex that differs in the neighbor simplex.

Inductive Hypothesis: Let $S = [\mathbf{v}_0 \dots \mathbf{v}_\ell \dots \mathbf{v}_d]^T$ be a simplex at level $\ell = |p| \bmod d$. Let $\ell^- = (\ell - 1) \bmod d$ and $\ell^* = \ell^- + 1$.

if (S is a 0-child) :

$$\begin{aligned} S^{(0)} &: [\mathbf{u} \ \mathbf{v}_1 \dots \mathbf{v}_d]^T, & \Pi^{(0)} &= \Pi \circ \Gamma_{\text{NEG},1} \\ S^{(i)} &: [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \ \mathbf{u} \ \mathbf{v}_{i+1} \dots \mathbf{v}_d]^T, & \Pi^{(i)} &= \Pi \circ \Gamma_{\text{SWP},i} \\ & \quad (0 < i < d) \end{aligned}$$

$$S^{(d)} : [\mathbf{v}_0 \dots \mathbf{v}_{\ell-1} \ \mathbf{u} \ \mathbf{v}_{\ell^*} \dots \mathbf{v}_{d-1}]^T, \quad \Pi^{(d)} = \Pi \circ \Gamma_{\text{RGT},\ell^-}$$

if (S is a 1-child) :

$$\begin{aligned} S^{(0)} &: [\mathbf{u} \ \mathbf{v}_1 \dots \mathbf{v}_d]^T, & \Pi^{(0)} &= \Pi \circ \Gamma_{\text{NEG},1} \\ S^{(\ell^*)} &: [\mathbf{v}_0 \dots \mathbf{v}_{\ell-1} \ \mathbf{v}_{\ell-+2} \dots \mathbf{v}_d \ \mathbf{u}]^T, & \Pi^{(\ell^*)} &= \Pi \circ \Gamma_{\text{LFT},\ell^-} \\ S^{(i)} &: [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \ \mathbf{u} \ \mathbf{v}_{i+1} \dots \mathbf{v}_d]^T, & \Pi^{(i)} &= \Pi \circ \Gamma_{\text{SWP},i} \\ & \quad (0 < i < d, i \neq \ell^*) \end{aligned}$$

$$S^{(d)} : [\mathbf{v}_0 \dots \mathbf{v}_{d-1} \ \mathbf{u}]^T \quad (d \neq \ell^*), \quad \Pi^{(d)} = \Pi \circ \Gamma_{\text{NSW},\ell}$$

Induction Basis: Neighbor rules hold for the $d!$ root simplices. Note that the *level* of a root simplex is 0, and the rules are the same whether the simplex is a 0-child, or a 1-child. Let S denote any root simplex, with *LPT code* $\Pi = [\pi_1 \dots \pi_i \ \pi_{i+1} \dots \pi_d]$.

For all root simplices, $S^{(0)} = \emptyset$, and $S^{(d)} = \emptyset$, that is, the 0^{th} and the d^{th} neighbors do not exist, since they are outside the *reference hypercube*. Other neighbors, $S^{(i)}$, $0 < i < d$, should be obtainable by swaps. Recall that the *base simplex* S_\emptyset and any root simplex S could be represented respectively as,

$$S_\emptyset = [\mathbf{y}_1 \dots \mathbf{y}_d], \quad \mathbf{y}_i = \begin{bmatrix} y_{i,0} \\ \dots \\ y_{i,i-1} \\ y_{i,i} \\ \dots \\ y_{i,d} \end{bmatrix} = \begin{bmatrix} -1 \\ \dots \\ -1 \\ 1 \\ \dots \\ 1 \end{bmatrix}$$

$$S = [\mathbf{y}'_1 \dots \mathbf{y}'_d], \quad \mathbf{y}'_j = \mathbf{y}_i \quad \text{iff} \quad \pi_i = j$$

If $\pi_i = j$, and $\pi_{i+1} = k$, then $\mathbf{y}'_j = \mathbf{y}_i$ and $\mathbf{y}'_k = \mathbf{y}_{i+1}$.

$$\Pi = [\pi_1 \dots \pi_{i-1} \ j \ k \ \pi_{i+2} \dots \pi_d]$$

Note that swapping columns \mathbf{y}'_j and \mathbf{y}'_k of S , will give us another valid root simplex, S' which differs from

S only in the i^{th} row, that is the i^{th} vertex. So, S' is basically the i^{th} neighbor of S , that is $S' = S^{(i)}$. Let Π' denote the signed permutation for S' . Then, $\Pi' = [\pi_1 \dots \pi_{i-1} \quad k \quad j \quad \pi_{i+2} \dots \pi_d]$. And so,

$$\Pi^{(i)} = \Pi' = \Pi \circ \Gamma_{\text{SWP},i}.$$

Induction Step: Let S be a simplex at level ℓ^- such that the inductive hypothesis holds. We will show that the inductive hypothesis holds for the two children of S . (Similar analysis can be used for the omitted cases.)

1. First, consider the 0-child of S , that is S_0 . Let ℓ denote the level of S_0 . Note that $\ell^- = (\ell - 1) \bmod d$. Let $\ell^* = \ell^- + 1$. Letting i denote the neighbor number, there are multiple cases to be distinguished.

(a) $i = 0$

If $0 < \ell^- \leq d - 1$,

$$\begin{aligned} S &= [\mathbf{v}_0 \quad \mathbf{v}_1 \dots \mathbf{v}_{\ell^-} \dots \mathbf{v}_d]^T \\ S^{(0)} &= [\mathbf{u} \quad \mathbf{v}_1 \dots \mathbf{v}_{\ell^-} \dots \mathbf{v}_d]^T \\ S_0 &= [\mathbf{v}_0 \quad \mathbf{v}_1 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \\ S_0^{(0)} &= [\mathbf{u} \quad \mathbf{v}_1 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \end{aligned}$$

Else if $\ell^- = 0$

$$\begin{aligned} S &= [\mathbf{v}_0 \quad \mathbf{v}_1 \dots \mathbf{v}_d]^T, \quad S_0 = [\mathbf{m} \quad \mathbf{v}_1 \dots \mathbf{v}_d]^T \\ S^{(0)} &= [\mathbf{u} \quad \mathbf{v}_1 \dots \mathbf{v}_d]^T, \quad S_0^{(0)} = [\mathbf{m}' \quad \mathbf{v}_1 \dots \mathbf{v}_d]^T \end{aligned}$$

In either case, $(S_0)^{(0)} = S_0^{(0)}$. Thus, $(\Pi_0)^{(0)} = \Pi_0^{(0)} = \Pi^{(0)} = \Pi \circ \Gamma_{\text{NEG},1} = \Pi_0 \circ \Gamma_{\text{NEG},1}$.

(b) $i = d$

By definition of bisection rules, d^{th} neighbor of S_0 is its sibling, that is S_1 , and

$$(\Pi_0)^{(d)} = \Pi_1 = \Pi_0 \circ \Gamma_{\text{RGT},\ell^-}.$$

(c) $0 < i < \ell^-$

$$\begin{aligned} S &= [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \quad \mathbf{v}_i \quad \mathbf{v}_{i+1} \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_d]^T \\ S^{(i)} &= [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \quad \mathbf{u} \quad \mathbf{v}_{i+1} \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_d]^T \\ S_0 &= [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \quad \mathbf{v}_i \quad \mathbf{v}_{i+1} \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \\ S_0^{(i)} &= [\mathbf{v}_0 \dots \mathbf{v}_{i-1} \quad \mathbf{u} \quad \mathbf{v}_{i+1} \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \end{aligned}$$

Then, $(S_0)^{(i)} = S_0^{(i)}$. Thus, $(\Pi_0)^{(i)} = \Pi_0^{(i)} = \Pi^{(i)} = \Pi \circ \Gamma_{\text{SWP},i} = \Pi_0 \circ \Gamma_{\text{SWP},i}$.

(d) $i = \ell^-, \ell^- \neq 0$

i. If S is a 0-child, (omitted, analogous to 1(c)).

ii. If S is a 1-child,

$$\begin{aligned} S &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^-} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \\ S^{(\ell^-)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d \quad \mathbf{u}]^T \\ S_0 &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \\ S_1^{(\ell^-)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m}' \quad \mathbf{v}_{\ell^*} \dots \mathbf{v}_d]^T \end{aligned}$$

Then, $(S_0)^{(\ell^-)} = S_1^{(\ell^-)}$.

$$\Pi_0 = \Pi = [\pi_1 \dots \pi_d]$$

$$\Pi^{(\ell^-)} = [\pi_1 \dots \pi_{\ell^- - 1} \quad \pi_{\ell^*} \dots \pi_d \quad -\pi_{\ell^-}]$$

$$\Pi_1^{(\ell^-)} = [\pi_1 \dots \pi_{\ell^- - 1} \quad \pi_{\ell^*} \quad \pi_{\ell^-} \quad \pi_{\ell^* + 1} \dots \pi_d]$$

Thus, $(\Pi_0)^{(\ell^-)} = \Pi_1^{(\ell^-)} = \Pi_0 \circ \Gamma_{\text{SWP},\ell^-}$.

(e) $\ell^- < i < d$, (omitted, analogous to 1(c)).

2. Next, consider the 1-child of S , that is S_1 . Let ℓ denote the level of S_1 . Note that $\ell^- = (\ell - 1) \bmod d$. Let $\ell^* = \ell^- + 1$. Letting i denote the neighbor number,

(a) $i = 0$

i. $\ell^- = 0$

$$\begin{aligned} S &= [\mathbf{v}_0 \quad \mathbf{v}_1 \dots \mathbf{v}_{d-1} \quad \mathbf{v}_d]^T \\ S^{(d)} &= [\mathbf{v}_0 \quad \mathbf{v}_1 \dots \mathbf{v}_{d-1} \quad \mathbf{u}]^T \\ S_1 &= [\mathbf{m} \quad \mathbf{v}_0 \dots \mathbf{v}_{d-1}]^T \\ S_1^{(d)} &= [\mathbf{m}' \quad \mathbf{v}_0 \dots \mathbf{v}_{d-1}]^T \end{aligned}$$

Then, $(S_1)^{(0)} = S_1^{(d)}$.

$$\Pi = [\pi_1 \dots \pi_d]$$

$$\Pi_1 = [-\pi_d \quad \pi_1 \dots \pi_{d-1}]$$

$$\Pi^{(d)} = [\pi_1 \dots \pi_{d-1} \quad -\pi_d]$$

$$\Pi_1^{(d)} = [\pi_d \quad \pi_1 \dots \pi_{d-1}]$$

Thus, $(\Pi_1)^{(0)} = \Pi_1^{(d)} = \Pi_1 \circ \Gamma_{\text{NEG},1}$.

ii. $\ell^- \neq 0$, (omitted).

(b) $0 < i < \ell^-$, (omitted).

(c) $i = \ell^-$

i. If S is a 0-child, (omitted).

ii. If S is a 1-child,

$$\begin{aligned} S &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-1} \quad \mathbf{v}_d]^T \\ S^{(d)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-1} \quad \mathbf{u}]^T \\ S_1 &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-1}]^T \\ S_1^{(d)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m}' \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-1}]^T \end{aligned}$$

Then, $(S_1)^{(\ell^-)} = S_1^{(d)}$.

$$\Pi = [\pi_1 \dots \pi_d],$$

$$\Pi_1 = [\pi_1 \dots \pi_{\ell^-} \quad -\pi_d \quad \pi_{\ell^*} \dots \pi_{d-1}]$$

$$\Pi^{(d)} = [\pi_1 \dots \pi_{\ell^- - 1} \quad -\pi_d \quad \pi_{\ell^*} \dots \pi_{d-1} \quad -\pi_{\ell^-}]$$

$$\Pi_1^{(d)} = [\pi_1 \dots \pi_{\ell^- - 1} \quad -\pi_d \quad \pi_{\ell^-} \dots \pi_{d-1}]$$

And so, $(\Pi_1)^{(\ell^-)} = \Pi_1^{(d)} = \Pi_1 \circ \Gamma_{\text{SWP},\ell^-}$.

(d) $i = \ell^*$, (omitted, analogous to 1(b)).

(e) $\ell^* < i \leq d$, $\ell \neq 0$

i. $\ell^* < i < d$, $\ell \neq 0$, (omitted, analogous to 2(e)ii)

ii. $i = d$, $\ell \neq 0$,

$$\begin{aligned} S &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-2} \quad \mathbf{v}_{d-1} \quad \mathbf{v}_d]^T \\ S^{(d-1)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-2} \quad \mathbf{u} \quad \mathbf{v}_d]^T \\ S_1 &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-2} \quad \mathbf{v}_{d-1}]^T \\ S_1^{(d-1)} &= [\mathbf{v}_0 \dots \mathbf{v}_{\ell^- - 1} \quad \mathbf{m} \quad \mathbf{v}_{\ell^-} \dots \mathbf{v}_{d-2} \quad \mathbf{u}]^T \end{aligned}$$

Then, $(S_1)^{(d)} = S_1^{(d-1)}$.

$$\Pi = [\pi_1 \dots \pi_{\ell^-} \dots \pi_{d-1} \quad \pi_d]$$

$$\Pi_1 = [\pi_1 \dots \pi_{\ell^-} \quad -\pi_d \quad \pi_{\ell^*} \dots \pi_{d-1}]$$

$$\Pi^{(d-1)} = [\pi_1 \dots \pi_{d-2} \quad \pi_d \quad \pi_{d-1}]$$

$$\Pi_1^{(d-1)} = [\pi_1 \dots \pi_{\ell^-} \quad -\pi_{d-1} \quad \pi_{\ell^*} \dots \pi_{d-2} \quad \pi_d]$$

Thus, $(\Pi_1)^{(d)} = \Pi_1^{(d-1)} = \Pi_1 \circ \Gamma_{\text{NSW},\ell^*} = \Pi_1 \circ \Gamma_{\text{NSW},\ell}$.