Dual Complexes of Cubical Subdivisions of $\mathbb{R}^n$

Herbert Edelsbrunner · Michael Kerber

Received: 16 November 2010 / Revised: 18 August 2011 / Accepted: 16 October 2011 /
Published online: 16 November 2011
© Springer Science+Business Media, LLC 2011

Abstract We use a distortion to define the dual complex of a cubical subdivision of $\mathbb{R}^n$ as an n -dimensional subcomplex of the nerve of the set of n -cubes. Motivated by the topological analysis of high-dimensional digital image data, we consider such subdivisions defined by generalizations of quad- and oct-trees to n dimensions. Assuming the subdivision is balanced, we show that mapping each vertex to the center of the corresponding n -cube gives a geometric realization of the dual complex in $\mathbb{R}^n$.

Keywords Simplicial complexes · (Hierarchical) cubical subdivisions · Counting · Distortion · Freudenthal triangulation · Geometric realization

1 Introduction

We are interested in cubical subdivisions of $\mathbb{R}^n$ as a generalization of the quad-tree and oct-tree data structures commonly used for 2- and 3-dimensional images [14, 15]. Thinking of an image as a discrete representation of a real-valued function, we view these trees as hierarchical representations and piecewise constant approximations of the same. The extension to $n \geq 4$ dimensions is motivated by the availability of high-resolution time-series of 3-dimensional images (e.g. Stock [16] observing the breaking of bone structure under pressure) and by the general quest to analyze multi-variate scientific data [8, 9]. While quad-trees and their higher-dimensional

H. Edelsbrunner · M. Kerber (✉)
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
e-mail: mkerber@ist.ac.at

H. Edelsbrunner
Departments of Computer Science and of Mathematics, Duke University, Durham, NC, USA

H. Edelsbrunner
Geomagic, Research Triangle Park, NC, USA

analogues are natural and efficient hierarchical representations of image data, a conversion to a simplicial representation facilitating a piecewise linear approximation is sometimes advantageous. An example is the construction of level sets, which requires a continuous function as input. In this work, we define and study such a conversion: the *dual complex* of a cubical subdivision. We build on Freudenthal's early work on triangulations of the n -dimensional cube [7]. The main results of this paper are as follows:

- I. We introduce a distortion of the integer grid in $\mathbb{R}^n$ to generalize the Freudenthal triangulation of the n -cube to the dual complex of a cubical subdivision of $\mathbb{R}^n$.
- II. We analyze the dual complex, giving tight bounds on its size and a detailed description of its local structure.
- III. We show that using the cube centers as the vertices of the dual complex of a balanced hierarchical cubical subdivision gives a geometric realization in $\mathbb{R}^n$.

The work reported in this paper is motivated by the desire to compute the persistent homology of n -dimensional images, generalizing the work of [2]. Indeed, from the dual complex of a cubical subdivision, we get an approximation of the image's persistent homology using standard reduction algorithms [6]. Alternatively, the same algorithms can be applied directly to the cubical grid [10]. Enhancements of the reduction algorithm based on collapses within the complex can be found in [5]. For images of moderate size, the persistent homology can be computed directly from the uniform cubical subdivision. While the size of this image representation may be large, its regular structure is amenable to implicit memory referencing, leading to implementations that are an order of magnitude faster than their counterparts for simplicial complexes of similar size [17]. In this context, it is important to distinguish between the *piecewise constant* approximations of a function furnished by cubical subdivisions, and the *piecewise linear* approximations provided by their dual complexes. The number of elements needed to achieve the same degree of approximation can be significantly smaller for the latter. Therefore, replacing a cubical subdivision by its dual complex can be an attractive option, in spite of the existence of algorithm that work directly on the primary structure.

Computing a simplicial refinement of a cubical subdivision is a well-studied problem in geometry processing, especially in two and three dimensions; see [18] for an extensive discussion with numerous references. All these approaches work by decomposing the cubes, with the constraint that the decompositions agree on the shared faces; see e.g. [4, Chap. 14]. In contrast, our dual complex of a cubical subdivision is not a refinement and instead consists of simplices connecting cube centers. As a consequence, it generally has fewer simplices than the refinements.

Outline Section 2 reviews the Freudenthal triangulation of the n -cube and counts its simplices. Section 3 explains the distortion and uses it to define the dual of a subdivision into unit cubes. Section 4 generalizes the construction to cubical subdivisions of nonuniform size. Section 5 introduces dual complexes and proves the geometric realization for balanced hierarchical subdivisions. Section 6 concludes the paper.

2 Freudenthal's Triangulation

In this section, we review the Freudenthal triangulation [7], also known as the Kuhn subdivision [11] of the n -dimensional cube.

The n -Cube The *unit n -cube* is the n -fold Cartesian product of the unit interval: $\mathbb{U}^n = [0, 1]^n \subseteq \mathbb{R}^n$. Picking $k \leq n$ of the intervals and either 0 or 1 from each of the remaining $n - k$ intervals, we get a k -face of $\mathbb{U}^n$, which is itself a k -dimensional cube. The number of k -faces is therefore

$$\bar{c}_k^n = \binom{n}{k} 2^{n-k}, \quad (1)$$

for all $0 \leq k \leq n$. To distinguish between different classes of faces, we write $\mathbf{0} = (0, 0, \dots, 0)$ and $\mathbf{1} = (1, 1, \dots, 1)$ for the extreme vertices in the diagonal direction, calling a face of $\mathbb{U}^n$ *anchored* at $\mathbf{0}$ (or $\mathbf{1}$) if it contains $\mathbf{0}$ (or $\mathbf{1}$) as one of its vertices. Some faces are anchored at $\mathbf{0}$, some are anchored at $\mathbf{1}$, and some are anchored at neither. Only one face of $\mathbb{U}^n$ is anchored at both, namely the n -cube itself, which is its only n -face. For each choice of k unit intervals, the only k -face anchored at $\mathbf{0}$ is the one for which the other $n - k$ coordinates are 0. Hence, the number of k -faces anchored at $\mathbf{0}$ is

$$\bar{a}_k^n = \binom{n}{k}, \quad (2)$$

for all $0 \leq k \leq n$. We are also interested in the silhouette of the n -cube, when viewed along the diagonal direction. For this reason, we introduce $\Delta : \mathbb{R}^n \rightarrow \mathbb{R}$ defined by mapping a point $x = (x_1, x_2, \dots, x_n)$ to $\Delta(x) = \sum_{i=1}^n x_i$. We refer to Δ as the *diagonal height function*, noting that $\Delta^{-1}(0)$ is the $(n - 1)$ -dimensional plane normal to the diagonal direction that pass through the origin, and $\Delta(x)$ is $\pm\sqrt{n}$ times the Euclidean distance from that $(n - 1)$ -plane. The orthogonal projection of the n -cube onto $\Delta^{-1}(0)$ is an $(n - 1)$ -dimensional convex polytope. This polytope has two decompositions into projections of $(n - 1)$ -cubes, generated by the $(n - 1)$ -faces of $\mathbb{U}^n$ anchored at $\mathbf{0}$ and by the $(n - 1)$ -faces anchored at $\mathbf{1}$. The *silhouette* of $\mathbb{U}^n$ consists of all points whose projection belongs to the boundary of that $(n - 1)$ -polytope. A face belongs to the silhouette iff it is neither anchored at $\mathbf{0}$ nor at $\mathbf{1}$. Indeed, each such face is shared by an $(n - 1)$ -face anchored at $\mathbf{0}$ and another anchored at $\mathbf{1}$. It is therefore easy to count them. Specifically, the number of k -faces in the silhouette of $\mathbb{U}^n$ is

$$\bar{s}_k^n = \binom{n}{k} (2^{n-k} - 2), \quad (3)$$

for all $0 \leq k \leq n - 1$. Since the silhouette is $(n - 2)$ -dimensional, the number of k -faces vanishes for $k = n - 1, n$. In Table 1, we give the number of faces, anchored faces, and silhouette faces for a few small values of n and k .

Table 1 From left to right in each entry: the number of k -faces of the n -cube, the number of k -faces anchored at **0** (or at **1**), and the number of k -faces in the silhouette. Zeros are omitted

	$k = 0$	1	2	3	4
$n = 1$	2, 1	1, 1			
2	4, 1, 1	4, 2	1, 1		
3	8, 1, 6	12, 3, 6	6, 3	1, 1	
4	16, 1, 14	32, 4, 24	24, 6, 12	8, 4	1, 1

Fig. 1 The Freudenthal triangulation of the 3-cube consists of six tetrahedra, arranged cyclically around the space-diagonal connecting **0** with **1**, and the faces of these tetrahedra

Chains We triangulate the n -cube using increasing sequences in a partial order of its vertices. Writing $\mathbf{i} = (i_1, i_2, \dots, i_n)$ and $\mathbf{j} = (j_1, j_2, \dots, j_n)$, with $i_k, j_k \in \{0, 1\}$ for all k , we say $\mathbf{i}$ *precedes* $\mathbf{j}$ if $i_k \leq j_k$ for all k . A *chain* is a sequence of distinct vertices in which each vertex precedes the next one in the partial order. Its *length* is the number of vertices. A chain is *maximal* if its length is $n + 1$. Each chain of length $k + 1$ defines a k -simplex, namely the convex hull of its $k + 1$ vertices. *Freudenthal's triangulation* of the n -cube, denoted by $\mathcal{F}^n = \mathcal{F}(\mathbb{U}^n)$, is the set of all simplices defined by chains [7]; see Fig. 1.

A maximal chain corresponds to a schedule of changing n 0's to n 1's, one coordinate at a time. It follows that there are $n!$ maximal chains, and similarly there are $n!$ n -simplices in $\mathcal{F}^n$. To count the k -simplices, we partition the set of n coordinate directions into $k + 2$ color classes, which we label from 0 to $k + 1$. Here we require that each color class between 1 and k contain at least one direction; the classes 0 and $k + 1$ may or may not contain directions. A maximal chain is *compatible* with a $(k + 2)$ -coloring if the coordinate directions that connect the vertices in sequence are ordered by color, from 0 to $k + 1$. Note that any two maximal chains compatible with the same $(k + 2)$ -coloring agree on the vertices that transition from one color to the next. We can therefore use the $(k + 2)$ -coloring to identify a unique k -simplex, namely the convex hull of the transition vertices, from the beginning of color 1 to the end of color k .

The number of $(k + 2)$ -colorings of the n coordinate directions is $(k + 2)^n$. To count the colorings that use each of the middle k colors at least once, we count all $(k + 2)$ -colorings and subtract the ones that do not use at least one of the middle colors. Note that $(k + 2 - i)^n$ of the $(k + 2)$ -colorings do not use some fixed subset of i colors. Inclusion-exclusion now implies that the number of k -simplices in the

Table 2 From left to right in each entry: the number of k -simplices in the Freudenthal triangulation of the n -cube, the number anchored at $\mathbf{0}$, and the number in the silhouette. Zeros are omitted

	$k = 0$	1	2	3	4
$n = 1$	2, 1	1, 1			
2	4, 1, 2	5, 3	2, 2		
3	8, 1, 6	19, 7, 6	18, 12	6, 6	
4	16, 1, 14	65, 15, 36	110, 50, 24	84, 60	24, 24

Freudenthal triangulation of the n -cube is

$$c_k^n = \sum_{i=0}^k (-1)^i \binom{k}{i} (k+2-i)^n, \quad (4)$$

for all $0 \leq k \leq n$. It is easy to see that this formula gives $c_0^n = 2^n$ but not quite as easy that it gives $c_n^n = n!$.

Anchors and Silhouettes A simplex is anchored at $\mathbf{0}$ iff color 0 is not used. We can therefore drop color 0 and compute the number of k -simplices in the Freudenthal triangulation of the n -cube that are anchored at $\mathbf{0}$ by counting $(k+1)$ -colorings as

$$a_k^n = \sum_{i=0}^k (-1)^i \binom{k}{i} (k+1-i)^n, \quad (5)$$

for all $0 \leq k \leq n$. If we now subtract the number of simplices anchored at $\mathbf{0}$ or at $\mathbf{1}$ from c_k^n , we get the number of k -simplices that triangulate the silhouette of the n -cube. We still need the number of k -simplices anchored at both, $\mathbf{0}$ and $\mathbf{1}$, which we get by counting k -colorings: $d_k^n = \sum_{i=0}^k (-1)^i \binom{k}{i} (k-i)^n$, for all $0 \leq k \leq n$. The number of k -simplices that triangulate the silhouette is therefore

$$s_k^n = c_k^n - 2a_k^n + d_k^n, \quad (6)$$

for all $0 \leq k \leq n$. Similar to the number of faces, we get $s_k^n = 0$ for $k = n-1, n$. We note that $s_k^n = d_{k+2}^n$ because the $(k+2)$ -colorings count the $(k+2)$ -simplices anchored at both $\mathbf{0}$ and $\mathbf{1}$, and each such $(k+2)$ -simplex has a unique k -face that is anchored at neither. In Table 2, we give the number of simplices, anchored simplices, and silhouette simplices for a few small values of n and k .

We note relations between the number of anchored simplices and the number of simplices in the silhouette, in the same and in one higher dimension. To express the relations without special cases, we set $s_{-1}^n = 1$ and $s_{-2}^n = 0$ for all dimensions n .

Anchor Formulas We have $a_k^n = s_{k-1}^n + s_{k-2}^n$ and $a_k^n = s_{k-1}^{n+1}/(k+1)$, for all $0 \leq k \leq n$.

Proof We employ straightforward algebraic manipulations to prove both relations. Using $\binom{k}{i} = \binom{k+1}{i} - \binom{k}{i-1}$, we get

$$\begin{aligned} a_k^n &= \sum_{i=0}^k (-1)^i \binom{k+1}{i} (k+1-i)^n \\ &\quad - \sum_{i=1}^k (-1)^i \binom{k}{i-1} (k+1-i)^n. \end{aligned}$$

Adding the vanishing term for $i = k+1$, we note that the first sum is d_{k+1}^n . Adding the vanishing term for $i = k+1$ and then transforming the index, we note that the second sum is $-d_k^n$. The first relation now follows from $s_{k-1}^n = d_{k+1}^n$ and $s_{k-2}^n = d_k^n$. Using $\binom{k}{i} = \frac{k+1-i}{k+1} \binom{k+1}{i}$, we get

$$a_k^n = \frac{k+1-i}{k+1} \sum_{i=0}^k (-1)^i \binom{k+1}{i} (k+1-i)^n.$$

Moving the factor $k+1-i$ into the sum and adding the vanishing term for $i = k+1$, we note that the sum is d_{k+1}^{n+1} . The second relation follows from $s_{k-1}^{n+1} = d_{k+1}^{n+1}$. $\square$

3 Uniform Cubical Subdivisions

The circumscribed $(n-1)$ -sphere of every n -simplex in $\mathcal{F}^n$ passes through the 2^n vertices of the unit n -cube. The Freudenthal triangulation is therefore a degenerate Delaunay triangulation. In this section, we study a distortion that selects $\mathcal{F}^n$ among all degenerate Delaunay triangulations.

Distortion in Diagonal Direction Write $\mathbb{Z}^n$ for the set of integer points in $\mathbb{R}^n$, and recall that the *Voronoi diagram* assigns to each point $\mathbf{i} \in \mathbb{Z}^n$ the cell of points $x \in \mathbb{R}^n$ for which $\mathbf{i}$ is a closest integer point. For $\mathbf{i} = (i_1, i_2, \dots, i_n)$, this cell is the Cartesian product of the intervals $[i_k - \frac{1}{2}, i_k + \frac{1}{2}]$, for $1 \leq k \leq n$, which is a translate of the unit n -cube. To remove common intersections of more than $n+1$ cells, we move the integer points by slightly compressing $\mathbb{Z}^n$ along the diagonal direction. Choosing $0 < \varepsilon < 1$, we map $\mathbf{i}$ to

$$\begin{aligned} T_\varepsilon \mathbf{i} &= \mathbf{i} - \varepsilon \frac{\Delta(\mathbf{i})}{n} \mathbf{1} \\ &= \left(i_1 - \varepsilon \frac{\Delta(\mathbf{i})}{n}, i_2 - \varepsilon \frac{\Delta(\mathbf{i})}{n}, \dots, i_n - \varepsilon \frac{\Delta(\mathbf{i})}{n} \right). \end{aligned}$$

Here, T_ε is the linear transformation defined by mapping the k th unit coordinate vector, $\mathbf{e}_k$, to $\mathbf{e}_k - \frac{\varepsilon}{n} \mathbf{1}$. It is the identity for $\varepsilon = 0$ and the orthogonal projection onto $\Delta^{-1}(0)$ for $\varepsilon = 1$. With this, we get Voronoi cells that are simple convex polyhedra, all of the same shape, namely combinatorially the same as a truncated n -cube;

Fig. 2 Sketch of the Voronoi cell of an integer point after the distortion in $\mathbb{R}^3$. It has the combinatorial structure of a cube after truncating two vertices and six edges

see Fig. 2 for the 3-dimensional case. As we will see shortly, the intersection of any $k + 1$ Voronoi cells is either empty or an $(n - k)$ -dimensional convex polytope, and which case it is does not depend on the particular value of $\varepsilon \in (0, 1)$. The intersection of $n + 2$ or more Voronoi cells is necessarily empty. We can therefore take the nerve of the set of Voronoi cells and get an n -dimensional simplicial complex: the Delaunay triangulation of the distorted set of integer points. We draw this complex in $\mathbb{R}^n$, using the (undistorted) integer points as vertices. In other words, we draw the complex as a degenerate Delaunay triangulation of the integer points, denoting it by $\mathcal{D}^n(\varepsilon) = \mathcal{D}_\varepsilon(\mathbb{Z}^n)$.

Triangulation We now formally prove that the nerve of the set of Voronoi cells gives an n -dimensional simplicial complex. Moreover, we show that $\mathcal{D}^n(\varepsilon)$ triangulates every integer translate of the unit n -cube by a copy of its Freudenthal triangulation.

Triangulation Theorem $\mathcal{D}^n(\varepsilon) = \mathcal{F}^n + \mathbb{Z}^n$, for every $0 < \varepsilon < 1$.

Proof We give the proof in two steps, simplifying by fixing ε and dropping it from the notation. The first step is geometric and shows that the claimed identity holds for the 1-skeleta of $\mathcal{D}^n$ and $\mathcal{F}^n$. The second step is combinatorial and shows that if we have the same edges, in $\mathcal{D}^n$ and $\mathcal{F}^n + \mathbb{Z}^n$, then we must also have the same higher-dimensional simplices. To prepare the two steps, we note that all Voronoi cells are integer translates of each other. Hence, $\mathcal{D}^n$ is invariant under integer translation. It therefore suffices to prove that $\mathcal{D}^n$ contains $\mathcal{F}^n$.

In the first step, we show that an edge connecting two integer points belongs to $\mathcal{D}^n$ iff it is an integer translate of an edge in $\mathcal{F}^n$. It is not difficult to see that every edge in $\mathcal{D}^n$ connects two vertices of an integer translate of the unit n -cube, so we may as well assume that both endpoints are vertices of $\mathbb{U}^n$. Writing V for its set of vertices, we observe that $\mathbb{U}^n$ is the convex hull of V . Since the distortion is a linear transformation, and linear transformations preserve convexity, $T\mathbb{U}^n$ is the convex hull of TV . Let S be the $(n - 1)$ -sphere that circumscribes $\mathbb{U}^n$. Its center is $(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$ and its radius is $\frac{1}{2}\sqrt{n}$. Recall that $\Delta^{-1}(\frac{n}{2})$ is the $(n - 1)$ -plane orthogonal to the diagonal that passes through the center of S . It intersects S in an $(n - 2)$ -sphere, $E = S \cap \Delta^{-1}(\frac{n}{2})$, which we refer to as the *equator* of S . The image of S under the linear transformation, TS , is an $(n - 1)$ -dimensional ellipsoid. It has one axis of length $(1 - \varepsilon)\sqrt{n}$, in the direction of the diagonal, and $n - 1$ axes of length $\sqrt{n}$, all axes of TE , which is a translate of the equator. Consider now a k -dimensional plane, P , and

the image of its intersection with the $(n - 1)$ -sphere: $T(P \cap S) = TP \cap TS$. Assume first that P passes through the center of S . Then $P \cap S$ is a $(k - 1)$ -sphere, and unless P is orthogonal to the diagonal direction, $P \cap E$ is a $(k - 2)$ -sphere, both with radius $\frac{1}{2}\sqrt{n}$. It follows that $TP \cap TS$ is a $(k - 1)$ -dimensional ellipsoid with one axis of length between $(1 - \varepsilon)\sqrt{n}$ and $\sqrt{n}$ and $k - 1$ axes of length $\sqrt{n}$. Indeed, the latter are axes of $T(P \cap E)$, which is a translate of $P \cap E$. The first axis is strictly shorter than $\sqrt{n}$ unless $P \subseteq \Delta^{-1}(\frac{n}{2})$. To understand the case in which P does not pass through the center of S , we note that parallel k -planes give rise to homothetic ellipsoids. The short axis of such an ellipsoid is always in the direction closest to the diagonal of $\mathbb{R}^n$, connecting the points with minimum and maximum diagonal height.

Consider now two vertices of $\mathbb{U}^n$ and let k be the smallest dimension such that both belong to a common k -face of $\mathbb{U}^n$, which we denote as $\mathbb{U}^k$. It has 2^{k-1} antipodal pairs of vertices, the chosen pair being one. The vertices of each pair differ from each other in precisely k coordinates. Hence, there is only one antipodal pair whose vertices are related to each other by the partial order, namely the pair u_0, u_1 in which u_0 has 0's and u_1 has 1's where they differ. By definition of the Freudenthal triangulation, this pair forms an edge in $\mathcal{F}^n$. To show that is also forms an edge in $\mathcal{D}^n$, we let P be the k -plane spanned by $\mathbb{U}^k$ and note that u_0 and u_1 are the orthogonal projections of $\mathbf{0}$ and $\mathbf{1}$ onto P . For reasons of symmetry, this implies that among the points of $P \cap S$, u_0 minimizes and u_1 maximizes the diagonal height. It follows that u_0, u_1 is the closest diagonal pair after distortion, implying that u_0, u_1 are the endpoints of an edge in $\mathcal{D}^n$. In summary, we proved that two vertices of $\mathbb{U}^n$ are connected by an edge in $\mathcal{D}^n$ iff they are related to each other in the partial order. Hence, the 1-skeleton of $\mathcal{D}^n$ is equal to the union of integer translates of the 1-skeleton of $\mathcal{F}^n$.

In the second step, we extend the result from edges to higher-dimensional simplices. Of course, a simplex can belong to $\mathcal{D}^n$ only if all its edges belong to $\mathcal{D}^n$. Restricting ourselves to the unit n -cube, $\mathbb{U}^n$, the vertices of a simplex in $\mathcal{D}^n$ thus form a chain in the partial order. Since $\mathcal{F}^n$ contains all such simplices, we just need to show that $\mathcal{D}^n$ also contains all such simplices. But if it does not, then it would be missing at least one of the n -simplices of $\mathcal{F}^n$, leaving a hole in the covering of $\mathbb{R}^n$ by the simplices in $\mathcal{D}^n$. This contradicts the Nerve Theorem, which states that $\mathcal{D}^n$ has the same homotopy type as the union of Voronoi cells, namely the homotopy type of $\mathbb{R}^n$; see e.g. [6, Sect. III.2]. $\square$

Implicit in the statement of the above theorem is that the triangulation does not depend on the particular choice of ε in the open unit interval. It is therefore convenient to drop the parameter from the notation and to write $\mathcal{D}^n = \mathcal{D}^n(\varepsilon)$ throughout the remainder of this paper.

Ratios of Limits of Ratios Now we know enough about $\mathcal{D}^n$ to count its simplices. Since there are infinitely many, we form unions of vertex stars and consider the ratio of the number of k -simplices over the number of vertices. Finally, we take the limit, letting the number of vertices go to infinity. Recall that each simplex in $\mathcal{D}^n$ has a unique lowest vertex with respect to the diagonal height function Δ , and that it belongs to the Freudenthal triangulation of the n -cube with the same lowest vertex. Hence, the limit of the ratio is the same as the number of k -simplices anchored at $\mathbf{0}$,

Table 3 The ratio of the number of k -simplices in $\mathcal{D}^n$ over the number of k -cubes in $\mathcal{V}^n$, up to one decimal position. The last column gives the ratio of the sums over all k : $\sum_k a_k^n / \sum_k \bar{a}_k^n$

	$k = 0$	1	2	3	4	5	
$n = 1$	1.0	1.0					1.0
2	1.0	1.5	2.0				1.5
3	1.0	2.3	4.0	6.0			3.2
4	1.0	3.7	8.3	15.0	24.0		9.3
5	1.0	6.2	18.0	39.0	72.0	120.0	33.8

Fig. 3 The projection of a level in $\mathcal{D}^3$ to the plane $\Delta^{-1}(0)$, and its triangulation

counted in (5). Summing this over all k , we get the limit ratio for the total number of simplices over the number of vertices as $\sum_{k=0}^n a_k^n$.

It is instructive to compare these numbers with the corresponding ratio limits for the subdivision of $\mathbb{R}^n$ into unit cubes, which we denote by $\mathcal{V}^n$. Each k -dimensional cube in $\mathcal{V}^n$ has a unique lowest vertex, at which it is anchored. The limit of the number of k -cubes over the number of n -cubes is therefore $\bar{a}_k^n = \binom{n}{k}$; see (2). In Table 3, we show the ratios of the ratio limits for small values of n and k . For example, for $k = n$, we have $n!$ n -simplices per n -cube and therefore a ratio of $n!$.

Levels We gain further insight into the structure of $\mathcal{D}^n$ by studying its relationship with $\mathcal{D}^{n+1}$. For this purpose, we consider the collection of n -faces of integer translates of the unit $(n+1)$ -cube in $\mathbb{R}^{n+1}$. Each such n -face has a unique lowest vertex in the diagonal height direction of $\mathbb{R}^{n+1}$. We define *level* ℓ as the faces whose lowest vertices have diagonal height ℓ . Projecting the level ℓ n -faces orthogonally onto $\Delta^{-1}(0)$, we get a subdivision of $\mathbb{R}^n$ by distorted n -cubes, which we denote as $\mathcal{L}_\ell^n$; see Fig. 3. Let $\mathcal{D}_\ell^n$ be the further subdivision of $\mathcal{L}_\ell^n$ into the simplices we get by projecting the Freudenthal triangulations of the n -faces. For $n \geq 2$, we have $\mathcal{L}_\ell^n \neq \mathcal{L}_{\ell+j}^n$ unless j is a multiple of $n+1$. In contrast, the triangulations are all the same.

Level Lemma $\mathcal{D}_\ell^n = \mathcal{D}_{\ell+j}^n$ for all integers ℓ and j .

Proof It suffices to show $\mathcal{D}_0^n = \mathcal{D}_1^n$. Since a level consists of n -cubes in $\mathbb{R}^{n+1}$, its vertices come on $n + 1$ different diagonal heights, namely $0, 1, \dots, n$ for level 0. Removing the integer points at height 0 and adding the ones at height $n + 1$, we get the vertices for level 1. But the integer points at heights 0 and $n + 1$ have the same projections in $\Delta^{-1}(0)$. This implies that $\mathcal{D}_0^n$ and $\mathcal{D}_1^n$ have the same vertices. It remains to show that they also have the same simplices of dimension larger than zero.

Consider a simplex of $\mathcal{D}_0^n$, and assume without loss of generality that it is the projection of a simplex in the Freudenthal triangulation of a lower n -face of $\mathbb{U}^{n+1}$. The vertices of that simplex have diagonal heights between 0 and n and they form a chain in the partial order in $\mathbb{R}^{n+1}$. If none of its vertices has height 0, this is also a chain in level 1, hence its projection also belongs to $\mathcal{D}_1^n$. However, if $\mathbf{0}$ is one of the vertices of the simplex, then we need to replace it by $\mathbf{1}$. The remaining vertices in the chain all succeed $\mathbf{0}$ and they all precede $\mathbf{1}$ in the partial order. Hence, we get a chain on level 1, which implies again that the projection of the simplex also belongs to $\mathcal{D}_1^n$, as required. $\square$

Links Suppose now that $\mathbf{i}'$ is the orthogonal projection onto $\Delta^{-1}(0)$ of the integer point $\mathbf{i}$ at height $\ell = \Delta(\mathbf{i})$ in $\mathbb{R}^{n+1}$. Hence, $\mathbf{i}'$ is a vertex of $\mathcal{L}_\ell^n$, and the distorted n -cubes that share $\mathbf{i}'$ are the projections of the $n + 1$ lower n -faces of $\mathbb{U}^{n+1} + \mathbf{i}$. The link of $\mathbf{i}'$ in $\mathcal{D}_\ell^n$ is therefore the projection of the triangulated silhouette of that $(n + 1)$ -cube. By the Level Lemma, every vertex in $\mathcal{D}_\ell^n$ is combinatorially the same as every other vertex. This implies that all links are integer translates of each other and of the projection of the triangulated silhouette of $\mathbb{U}^{n+1}$. It is now not difficult to prove that a similar statement holds for the degenerate Delaunay triangulation $\mathcal{D}^n$ in $\mathbb{R}^n$.

Link Lemma *The links of the vertices in $\mathcal{D}^n$ are integer translates of each other, and they are all isomorphic to the triangulated silhouette of the unit $(n + 1)$ -cube.*

Proof The n -dimensional simplicial complexes $\mathcal{D}_\ell^n$ in $\Delta^{-1}(0)$ and $\mathcal{D}^n$ in $\mathbb{R}^n$ are geometrically different but combinatorially the same. Specifically, $\mathcal{D}^n$ is the (non-orthogonal) diagonal projection of a level in $\mathcal{D}^{n+1}$ onto the n -dimensional plane spanned by the first n coordinate axes. Hence, we get $\mathcal{D}_\ell^n$ as the image of $\mathcal{D}^n$ under the linear transformation T_ε , with $\varepsilon = 1 - 1/\sqrt{n + 1}$. This implies that $\mathcal{D}^n$ and $\mathcal{D}_\ell^n$ are isomorphic, so the links of their vertices are isomorphic. The second part of the claim follows because the vertex links in $\mathcal{D}_\ell^n$ are isomorphic to the triangulated silhouette of $\mathbb{U}^{n+1}$, by construction. $\square$

Since all vertex links in $\mathcal{D}^n$ are isomorphic to the triangulated silhouette of the $(n + 1)$ -cube, we can use the results of Sect. 2 to count their simplices. Specifically, the link of a vertex in $\mathcal{D}^n$ has s_k^{n+1} k -simplices, for $0 \leq k \leq n - 1$. It follows that the star of a vertex in $\mathcal{D}^n$ has s_{k-1}^{n+1} k -simplices. Since each k -simplex belongs to $k + 1$ vertex stars, the ratio of the number of k -simplices over the number of vertices is $s_{k-1}^{n+1}/(k + 1)$. By the second Anchor Formula, this is indeed equal to a_k^n .

Fig. 4 Left: a piece of a cubical subdivision of the plane with overlaid piece of the dual complex. Right: the fractally distorted images of the squares

4 Non-uniform Cubical Subdivisions

In this section, we extend the results from uniform to non-uniform cubical subdivisions, focusing on generalizations of quad- and oct-trees to hierarchical subdivisions of $\mathbb{R}^n$.

Cubical Subdivisions Recall the setting in Sect. 3, where we begin with the subdivision of $\mathbb{R}^n$ into unit n -cubes centered at the integer points. We relax the size requirement and consider subdivisions of $\mathbb{R}^n$ into n -cubes that are unions of these unit n -cubes. To avoid the otherwise easy confusion between n -cubes and unit n -cubes, we will refer to the former as *cells*.

Definition A *cubical subdivision* of $\mathbb{R}^n$ is a collection, $\mathcal{C}$, of n -dimensional cubical cells with disjoint interiors that cover $\mathbb{R}^n$, for which we require that each unit n -cube centered at a point in $\mathbb{Z}^n$ is contained in a cell in $\mathcal{C}$.

See Fig. 4 for a 2-dimensional example. By definition, each cell $C \in \mathcal{C}$ with edges of length ℓ is the union of ℓ^n unit n -cubes, $C = U_1 \cup U_2 \cup \dots \cup U_{\ell^n}$. Each U_j is the Voronoi cell of an integer point $\mathbf{i} \in \mathbb{Z}^n$, and corresponds to a distorted truncated cube, $U_j(\varepsilon)$, the Voronoi cell of the integer point after distortion, $T_\varepsilon \mathbf{i} \in T_\varepsilon \mathbb{Z}^n$. We call $C(\varepsilon) = U_1(\varepsilon) \cup U_2(\varepsilon) \cup \dots \cup U_{\ell^n}(\varepsilon)$ a *fractally distorted cell*. Note that $C(\varepsilon)$ is different from $T_\varepsilon C$, as can be seen in Fig. 4. Since the $U_i(\varepsilon)$ depend on ε , we get a 1-parameter family of fractally distorted cells $C(\varepsilon)$ for each $C \in \mathcal{C}$. Assuming $\ell \geq 2$, $C(\varepsilon)$ is not convex for any positive ε but has a convex limit, at $\varepsilon = 0$.

Distorted Intersections Let now $C_0, C_1, \dots, C_k$ be cells in a cubical subdivision, $F = \bigcap_{i=0}^k C_i$ their common intersection, and $F(\varepsilon) = \bigcap_{i=0}^k C_i(\varepsilon)$ the common intersection after distortion. Since the C_i are convex, F is either empty or convex. In contrast, $F(\varepsilon)$ is not necessarily convex. Furthermore, $F = \emptyset$ implies $F(\varepsilon) = \emptyset$, but not the other way round. To describe the relationship between a face before and after distortion, we consider the limit of $F(\varepsilon)$, for ε going to 0. It consists of all points x for which there are points $x(\varepsilon) \in F(\varepsilon)$ such that $x = \lim_{\varepsilon \rightarrow 0} x(\varepsilon)$. If the C_i are unit

n -cubes, then the limit of a non-empty $F(\varepsilon)$ is equal to F . More generally, all points x in the limit of $F(\varepsilon)$ belong to F , but there can be points $y \in F$ that are not in the limit of $F(\varepsilon)$. We now prove that such points y exist only if $F(\varepsilon) = \emptyset$.

Limit Lemma *If $F(\varepsilon) \neq \emptyset$ then $\lim_{\varepsilon \rightarrow 0} F(\varepsilon) = F$.*

Proof We assume $F(\varepsilon) \neq \emptyset$ and note that $\lim_{\varepsilon \rightarrow 0} F(\varepsilon) \subseteq F$. We prove equality indirectly, assuming there is a point $y \in F$ not in the limit of $F(\varepsilon)$. The interiors of the unit n -cubes and of their faces partition each C_i and therefore also F . Hence, there is a unique unit cube that contains y in its interior, and we suppose its dimension is maximal, that is, equal to $\ell = \dim F$. Let L be the ℓ -plane that contains this unit ℓ -cube, and let $U_0, U_1, \dots, U_k$ be a selection of unit n -cubes with $y \in U_i \subseteq C_i$, for $0 \leq i \leq k$. Let N be the $(n - \ell)$ -plane orthogonal to L that passes through the centers of the U_i . We may assume that N is defined by $x_{n-\ell+1} = x_{n-\ell+2} = \dots = x_n = 0$. The centers of the U_i do not form a chain, else y would be in the limit of $\bigcap_{i=0}^k U_i(\varepsilon) \subseteq F(\varepsilon)$. It follows that $\bigcap_{i=0}^k U_i(\varepsilon) = \emptyset$, for $\varepsilon > 0$. We need to prove that the same is true for every other selection of unit n -cubes $V_0, V_1, \dots, V_k$ with $V_i \subseteq C_i$, for $0 \leq i \leq k$. Note that we do not require that y belongs to the common intersection of the V_i . To get a contradiction, we assume the centers of the V_i form a chain. Define the *rectangular hull* of V_i and U_i as the collection of unit cubes W_i such that

$$\min\{u_{ij}, v_{ij}\} \leq w_{ij} \leq \max\{u_{ij}, v_{ij}\}$$

for each $1 \leq j \leq n$, where u_i is the center of the unit n -cube U_i , u_{ij} is its j th coordinate, and similar for v_i, v_{ij} and w_i, w_{ij} . Clearly, all W_i in the rectangular hull of V_i and U_i belong to C_i . Let V'_i be the unit n -cube whose center, v'_i , is the orthogonal projection of v_i onto N . In other words,

$$v'_{ij} = \begin{cases} v_{ij} & \text{for } 1 \leq j \leq n - \ell, \\ u_{ij} & \text{for } n - \ell < j \leq n. \end{cases}$$

Since V'_i belongs to the rectangular hull of V_i and U_i , it also belongs to C_i . It follows that the V'_i are $k + 1$ distinct unit n -cubes. But then, the v'_i inherit the property of forming a chain from the v_i . We have $y \in \bigcap_{i=0}^k V'_i$, since the v'_i all lie in N , which contradicts the assumption that y does not belong to the limit of $F(\varepsilon)$. Hence, the v'_i cannot form a chain, and neither can the v_i . It follows that $\lim_{\varepsilon \rightarrow 0} F(\varepsilon) = F$ whenever $F(\varepsilon) \neq \emptyset$, as claimed. $\square$

The contrapositive form of the Limit Lemma is perhaps a more vivid description of how a cubical subdivision relates to its fractally distorted image: if $F \neq \lim_{\varepsilon \rightarrow 0} F(\varepsilon)$ then $F(\varepsilon) = \emptyset$ for $\varepsilon > 0$. In particular, if the dimension of F exceeds $n - k$ then $F(\varepsilon) = \emptyset$.

Face Structure After distortion, the unit n -cubes form a simple cell complex. It follows that the non-empty intersection of $k + 1$ distorted unit n -cubes is necessarily $(n - k)$ -dimensional. Hence, $F(\varepsilon) = \bigcap_{i=0}^k C_i(\varepsilon)$ is either empty or $(n - k)$ -dimensional. In the latter case, it is not difficult to show that $F(\varepsilon)$ is an $(n - k)$ -dimensional manifold with boundary, for $\varepsilon > 0$. In the limit, for $\varepsilon = 0$, the common

Fig. 5 *Left:* the regular subdivision of F into unit ℓ -cubes, for $\ell = 2$. *Right:* the corresponding distorted ℓ -cubes with filled gaps between them

intersection is convex and therefore contractible. It is therefore plausible that $F(\varepsilon)$ is contractible also for $\varepsilon > 0$. This is implied by the following result.

Fractal Distortion Lemma *For every $0 \leq k \leq n$, the common intersection of the fractally distorted images of $k+1$ cells in a cubical subdivision of $\mathbb{R}^n$ is either empty or an $(n-k)$ -ball.*

Proof We give an explicit construction of $F(\varepsilon)$. Supposing $F(\varepsilon) \neq \emptyset$, we can find unit n -cubes $U_0, U_1, \dots, U_k$, with $U_i \subseteq C_i$ for each i , whose centers form a chain of length $k+1$. Here, we choose the indices so their ordering is consistent with the ordering of the centers along the chain. For each pair $0 \leq i < i' \leq k$, there is at least one coordinate direction, j , for which a normal $(n-1)$ -plane separates C_i from $C_{i'}$. We call j a *separating* coordinate direction for C_i and $C_{i'}$. The separating directions for C_0 and C_1 are distinct from those for C_1 and C_2 , and so on. Letting S be the collection of separating coordinate directions, we therefore have $|S| \geq k$. Let T be the complementary collection of non-separating coordinate directions, and note that $\dim F = n - |S| = |T|$, with $F = \bigcap_{i=1}^k C_i$, as before. Writing $\ell = |T|$, we know that F is an ℓ -dimensional rectangular box. For each unit ℓ -cube in its subdivision, we have a chain in which the first vertex and the last vertex differ in $n - \ell$ coordinates. Equivalently, their unit n -cubes have $n - \ell$ separating directions. The corresponding $k+1$ distorted unit n -cubes intersect in an $(n-k)$ -dimensional face whose limit, for $\varepsilon = 0$, is ℓ -dimensional. We project these $(n-k)$ -dimensional faces into an $(n-k)$ -plane, which we choose so that the images of the $(n-k)$ -faces are disjoint, as in Fig. 5. To construct this $(n-k)$ -plane, we select k coordinate directions, one each separating C_{i-1} and C_i , for $1 \leq i \leq k$. Finally, we take the distorted images of these directions and get the $(n-k)$ -plane as the intersection of the $(n-1)$ -planes normal to the distorted directions.

In the last step of our proof, we construct the faces that fill the gaps between the projections of the $(n-k)$ -faces whose limits are the unit ℓ -cubes decomposing F . These faces can be enumerated by moving the vertices in a chain one by one in a non-separating coordinate direction in such a way that the chain remains a chain. In other words, we use chains in which some of the directions in T separate the corresponding unit n -cubes. Letting the number of additional separating directions be $m \leq \ell$, the chain corresponds to an $(n-k)$ -face whose limit is $(\ell-m)$ -dimensional. Using all subsets of T and, for each subset, all chains for which the directions in the subset separate, we fill all gaps between the distorted ℓ -cubes. We may even get more, namely an incomplete extra layer of faces around the configuration of $(n-k)$ -faces whose limits are the unit ℓ -cubes decomposing F . In any case, the collection

of $(n - k)$ -faces forms an $(n - k)$ -dimensional ball whose limit, for $\varepsilon = 0$, is an ℓ -dimensional rectangular box. $\square$

Hierarchical Cubical Subdivisions We are interested in cubical subdivisions that arise from a hierarchical decomposition of $\mathbb{R}^n$, generalizing quad-trees in $\mathbb{R}^2$ and oct-trees in $\mathbb{R}^3$. To define them, we limit the set of available cells to a *basis* $\mathcal{B}$ of n -dimensional cubes B for which there are integers $\ell \geq 0$ and $m_1, m_2, \dots, m_n$ such that B is the union of the unit n -cubes centered at the integer points $(i_1, i_2, \dots, i_n)$ with $2^\ell m_k + 1 \leq i_k \leq 2^\ell (m_k + 1)$, for each $1 \leq k \leq n$. We call 2^ℓ the *size* of B . Taking all cubes of size 2^ℓ gives a uniform cubical subdivision of $\mathbb{R}^n$. Hence, we can think of $\mathcal{B}$ as a hierarchy of uniform subdivisions in which the number of cubes grows exponentially from one level to the next.

Definition A *hierarchical cubical subdivision* of $\mathbb{R}^n$ is a cubical subdivision $\mathcal{C} \subseteq \mathcal{B}$. Its *closure*, $\overline{\mathcal{C}}$, consists of all cubes in $\mathcal{B}$ that contain cubes in $\mathcal{C}$, and its *interior* is the closure minus the subdivision itself, $\mathcal{C}^\circ = \overline{\mathcal{C}} - \mathcal{C}$.

Every hierarchical cubical subdivision has a unique closure and a unique interior. Conversely, the closure determines the subdivision, and so does the interior. A *refinement* of $\mathcal{C}$ is a hierarchical cubical subdivision whose closure contains $\overline{\mathcal{C}}$. While hierarchical cubical subdivisions are necessarily infinite, we can extract finite pieces. Specifically, for each n -cube $B \in \overline{\mathcal{C}}$, we define $\mathcal{C}(B) = \{C \in \mathcal{C} \mid C \subseteq B\}$, referring to it as a *finite* hierarchical cubical subdivision. See Fig. 4 for an example in the plane. Accordingly, the closure and interior of $\mathcal{C}(B)$ are the subsets of cells in $\overline{\mathcal{C}}$ and $\mathcal{C}^\circ$ that are contained in B . In the finite case, the sizes of a subdivision, its closure, and its interior are tightly coupled:

$$\begin{aligned} |\overline{\mathcal{C}}(B)| &= |\mathcal{C}(B)| + |\mathcal{C}^\circ(B)| \\ &= 2^n |\mathcal{C}^\circ(B)| + 1. \end{aligned} \quad (7)$$

It should be clear that we can think of $\overline{\mathcal{C}}(B)$ as a tree in the computer science sense. Its cells are the *nodes*, distinguishing between the *internal nodes* in $\mathcal{C}^\circ(B)$ and the *external nodes* in $\mathcal{C}(B)$. The *children* of a node are the cells of half the size contained in it, and the *parent* is the cell of twice the size that contains it. Other than the *root* of the tree, which is B , every node has exactly one parent, every internal node has 2^n children, and every external node has no child.

Balancing We refer to cells whose fractally distorted images have a non-empty intersection as *neighbors*. Generalizing [3], we call a hierarchical cubical subdivision of $\mathbb{R}^n$ *balanced* if any two neighboring cells are either of the same size or one is twice the size of the other. For example, the quad-tree subdivision in Fig. 4 is not balanced as it has neighboring squares whose sizes differ by a factor of four. It is, however, easy to make it balanced, namely by subdividing the upper left square into four. It is not difficult to see that every hierarchical cubical subdivision has a unique smallest balanced refinement, which is obtained by greedily subdividing cells that violate the balancing condition [12]. To compare $\mathcal{C}$ with this refinement, we adapt a result of Moore [12]; see [4, Theorem 14.4] for the special case of $n = 2$.

Balancing Lemma *Let $\mathcal{C}$ be a hierarchical cubical subdivision of $\mathbb{R}^n$ and $\mathcal{R}_{\min}$ its smallest balanced refinement. Then $|\mathcal{R}_{\min}(B)| \leq 3^n |\mathcal{C}(B)|$ for every cell $B \in \overline{\mathcal{C}}$.*

Proof Call two cells in a subdivision *adjacent* if they have a non-empty intersection, and note that any two neighboring cells are adjacent but not the other way round. We call the subdivision *strongly balanced* if any two adjacent cells differ in size by at most a factor of two. Let $\mathcal{R}$ be the smallest strongly balanced refinement of $\mathcal{C}$. Since strong balance implies balance, it suffices to show that $|\mathcal{R}(B)| \leq 3^n |\mathcal{C}(B)|$ for every $B \in \overline{\mathcal{C}}$.

We use the language of trees to explain the proof. The key insight is the following: if an external node is split during the refinement process, then $\mathcal{C}^\circ$ contains a node adjacent to and of the same size as the split node; see [12] or [4, Theorem 14.4] for more details. It follows that every split in the process can be charged to a node of $\mathcal{C}^\circ$. Now, each internal node is charged at most $3^n - 1$ times because this is the number of adjacent nodes of same size. With every split, the number of external nodes increases by $2^n - 1$, so we get

$$\begin{aligned} |\mathcal{R}(B)| &\leq |\mathcal{C}(B)| + (2^n - 1)(3^n - 1)|\mathcal{C}^\circ(B)| \\ &\leq |\mathcal{C}(B)| + (3^n - 1)(|\mathcal{C}(B)| - 1), \end{aligned}$$

because $|\mathcal{C}(B)| - 1 = (2^n - 1)|\mathcal{C}^\circ(B)|$, which we get from (7). The claimed inequality follows. $\square$

5 Dual Complexes

In this section, we introduce the main new concept of this paper, namely the dual complex of a non-uniform cubical subdivision. It is not necessarily a Delaunay triangulation, so we have to worry about embedding it.

Triangulation Similar to the uniform case, we need the distortion to control the explosion in dimension we otherwise get by taking the nerve of a collection of cubes.

Definition The *dual complex* of a cubical subdivision $\mathcal{C}$ of $\mathbb{R}^n$ is the system of subsets $\mathcal{K}^n = \mathcal{K}(\mathcal{C})$ that contains $\alpha \subseteq \mathcal{C}$ if the fractually distorted images of the cells in α have a non-empty common intersection.

We extend this notion by calling the full subcomplex of $\mathcal{K}(\mathcal{C})$ defined by a subset of $\mathcal{C}$ the *dual complex* of the subset. Observe that the definition of the dual complex is independent of the particular choice of the parameter $\varepsilon \in (0, 1)$ that controls the distortion. Moreover, the dual complex can be computed purely combinatorially, without constructing distorted cells, as in [2]. We put $\mathcal{K}^n$ into $\mathbb{R}^n$ by mapping each cell to its (undistorted) center and drawing each subset of cells as the convex hull of their centers. This does not necessarily give a simplicial complex, in which any two simplices are either disjoint or intersect in a common face. However, we will identify an important class of cubical subdivisions for which this drawing of $\mathcal{K}^n$ is a geometric realization in $\mathbb{R}^n$.

Fig. 6 Cutting the *middle square* into two creates one new vertex and three new edges

Ratio Bounds Before addressing the question of geometric realization, we give an upper bound on the number of simplices in a dual complex. Recall that $\mathcal{K}^n = \mathcal{D}^n$ if all n -cubes are of unit size. As shown at the end of Sect. 3, in this case the ratio of the number of k -simplices over the number of vertices is a_k^n . We now show that this is the largest ratio we can get.

Size Lemma *The number of k -simplices over the number of vertices in the dual complex of a hierarchical cubical subdivision of $\mathbb{R}^n$ is at most a_k^n .*

Proof Our argument works by stepwise refinement of the subdivision $\mathcal{C}$ until we arrive at $\mathcal{V}^n$, in which all cells are unit n -cubes. We already have a good understanding of $\mathcal{D}^n = \mathcal{K}(\mathcal{V}^n)$. Specifically, the ratio of the number of k -simplices over the number of vertices in $\mathcal{D}^n$ is a_k^n ; see Sect. 3. We express this by saying that the average number of k -simplices per vertex is a_k^n . We will prove that each refinement step adds one vertex and at least a_k^n k -simplices. Since the average is a_k^n at the end, for $\mathcal{V}^n$, it cannot be more than a_k^n at the beginning, for $\mathcal{C}$.

We refine $\mathcal{C}$ by subdividing its cells in the order of non-increasing size, using $2^n - 1$ straight cuts to subdivide a cell into 2^n cells of half the size. We do these cuts in sequence but not consecutively, as we now explain. When we cut a cell, we get two rectangular boxes, each with *long* sides of the same length as the edges of the cell, and a *short* side of half the length. In general, we get boxes with $n - k$ long and k short sides, where k is anywhere between 0 and n . We order the cuts such that the short sides are parallel to the first k coordinate directions and the long sides are parallel to the last $n - k$ coordinate directions. To compare, we say a box B is *larger* than another box if the long sides of B are longer, or the long sides of the two boxes have equal length but B has more long sides. Finally, we refine $\mathcal{C}$ by cutting the boxes in the order of non-increasing size.

Let now B be a largest box and k its number of short edges. Because of the order of the cuts, the neighbors of B are smaller than or of the same size as B . We cut B in half, with an $(n - 1)$ -plane normal to the $(k + 1)$ -st coordinate direction. Cutting the box corresponds to splitting the corresponding vertex in the dual complex; see Fig. 6. A new edge connecting the two copies of the split vertex appears. The link of this edge is a triangulation of the $(n - 2)$ -sphere. We denote this link by L , observing that it is a subcomplex of the link of the vertex before the split. If all neighbors of B are of the same size as B , then L is isomorphic to a vertex link in $\mathcal{D}^{n-1}$; see the remark after the Link Lemma in Sect. 3. In this case, L has s_k^n k -simplices. If some of the neighbors of B are smaller, then the number of k -simplices in the link is at

Fig. 7 Three cubes in $\mathbb{R}^3$ whose centers span a triangle that is not contained in the union of the three cubes

least s_k^n . The split doubles the set of simplices connecting the vertex with simplices in L , and it triangulates the space in between. In other words, for each k -simplex in L , we get an additional $(k+1)$ -simplex by doubling and an additional $(k+2)$ -simplex by filling. Hence, the number of new k -simplices that appear as a result of the split is at least $s_{k-1}^n + s_{k-2}^n$. The result follows because this sum is equal to a_k^n by the first Anchor Formula in Sect. 2. $\square$

By the Size Lemma, a_k^n is the largest possible ratio between the number of k -simplices and vertices for non-uniform cubical subdivisions. On the other hand, the ratio between the number of k -cubes and n -cubes cannot be smaller than $\binom{n}{k}$. Therefore, Table 3 gives upper bounds for the number of k -simplices in the dual complex over the number of k -cubes in the cubical subdivision.

Counterexample to Geometric Realization We are now ready to address the question of geometric realization. For dimension $n = 2$, it is fairly easy to prove that the dual complex of a cubical subdivision is geometrically realized in $\mathbb{R}^2$. The key insight is that every edge of $\mathcal{K}^2$ is contained in the union of the two squares that define it; compare with Fig. 4. While this property generalizes to $\mathbb{R}^n$, it no longer implies the geometric realization of the dual complex. Following [2], we now describe a counterexample in three dimensions.

We begin with two cubes, A and B , that share a common edge of length 8. To this, we add a cube C of size 2 such that one of its edges overlaps with the last quarter of the shared edge of A and B ; see Fig. 7. The line segment connecting the centers of A and B passes through the midpoint of the shared edge. This midpoint lies outside C , and the center of C lies outside $A \cup B$. The line segment connecting the midpoint and this center belongs to the triangle spanned by the three centers but it is not contained in $A \cup B \cup C$. This implies that the triangle lies partially outside the three cubes. Now we just need to place a unit cube on top of C so it touches both A and B . Its center lies on the triangle and thus forms an improper intersection.

The configuration in Fig. 7 is part of a hierarchical cubical subdivision of $\mathbb{R}^3$. Note, however, that this subdivision is not balanced. In the remainder of this section, we show that balance prohibits improper intersections between simplices in the dual complex in all positive dimensions.

Fig. 8 Seed configuration of a tetrahedron in the dual complex of a cubical subdivision of $\mathbb{R}^3$. The *white dots* are the centers of the unit cubes in the seed configuration, and the *black dots* are the centers of the corresponding cubes of twice the size

Seed Configurations Let now $\mathcal{C}$ be a hierarchical cubical subdivision of $\mathbb{R}^n$, and let $C_0, C_1, \dots, C_n$ be cells in $\mathcal{C}$ forming an n -simplex in $\mathcal{K}^n = \mathcal{K}(\mathcal{C})$. By the Fractal Distortion Lemma, the corresponding fractally distorted cells meet in a single common point, which we denote as $T_\varepsilon x$. The coordinates of the corresponding undistorted point, x , are odd multiples of $\frac{1}{2}$. The point $T_\varepsilon x$ is also common to the distorted images of $n + 1$ unit n -cubes, one in each $C_k(\varepsilon)$. In other words, there is a unique collection of unit n -cubes $U_k \subseteq C_k$, for $0 \leq k \leq n$, such that

$$T_\varepsilon x = \bigcap_{k=0}^n C_k(\varepsilon) = \bigcap_{k=0}^n U_k(\varepsilon);$$

see Fig. 8. Writing u_k for the center of U_k , for each k , we call $u_0, u_1, \dots, u_n$ the *seed configuration* of the n -simplex. To study this configuration, we may assume that the u_k are vertices of $\mathbb{U}^n = [0, 1]^n$. Writing u_{kj} for the j th coordinate of u_k , we can make this more specific by assuming $u_{kj} = 1$ if $j \leq k$ and $u_{kj} = 0$ if $k < j$. The common point of the U_k is then $x = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$, the center of $\mathbb{U}^n$.

Two orderings of the vertices of an n -simplex belong to the same *orientation* if they differ by an even number of transpositions. Writing the vertices as the rows of a matrix, in the sequence of their ordering, and adding a column of 1's on the left, we can use the sign of the determinant to distinguish between the two orientations. For example, for the ordering $u_0, u_1, \dots, u_n$ we get

$$\det \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ 1 & 1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 1 \end{bmatrix} = 1, \quad (8)$$

and we say this ordered n -simplex has *positive orientation*. The determinant is also $n!$ times the signed n -dimensional volume of the n -simplex. Since the volume is a continuous function of the $n + 1$ points, we can move the points around and be sure the determinant does not change its sign, unless the points pass through a configuration in which they are affinely dependent. Because of this property, it is possible to compare the orientations of different n -simplices, as we will do extensively below.

Orientation In a geometrically realized dual complex, all n -simplices have the same orientation as their seed configurations. We now prove that dual complexes of balanced subdivisions have this property.

Orientation Lemma *Every n -simplex in the dual complex of a balanced hierarchical cubical subdivision of $\mathbb{R}^n$ has the same orientation as its seed configuration.*

Proof Let $C_0, C_1, \dots, C_n$ be a sequence of n -dimensional cubical cells in the balanced hierarchical cubical subdivision, assume they define an n -simplex in the dual complex, and let $U_0, U_1, \dots, U_n$ be the corresponding sequence of unit n -cubes in the seed configuration. We write c_k for the center of C_k and c_{kj} for its j th coordinate. It is convenient to assume that the seed configuration has the special form described above. Since the C_k come in at most two sizes, we may assume that either $C_k = U_k$ or C_k is twice the size of U_k . In the latter case, c_k is a vertex of U_k , and we have $|c_{kj} - u_{kj}| = \frac{1}{2}$ for all j . Assuming $k \neq \ell$ are indices with $c_k \neq u_k$ and $c_\ell \neq u_\ell$, the difference between the coordinates of their centers is

$$c_{kj} - c_{\ell j} \in \{-2, 0, 2\}, \quad (9)$$

for each $1 \leq j \leq n$. The difference is a multiple of 2 because C_k and C_ℓ are part of a hierarchical subdivision, and it cannot be larger than 2 because they are corners of neighboring unit n -cubes.

A particular choice for the center of C_k is $c_k = 2u_k - x$ with $x = (\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2})$; see the black dots in Fig. 8. Here, the coordinate vector of c_k consists of k leading $\frac{3}{2}$'s and $n - k$ trailing $-\frac{1}{2}$'s. We consider the case in which $c_k = 2u_k - x$, for some indices k , and $c_k = u_k$, for others. We claim that the orientation of the n -simplex is still positive. To see this, we consider again the matrix of vertex coordinates. The k th row is either the same as in (8) or different in the way described above. Let m be the smallest index for which $c_m \neq u_m$. We subtract row m from each row $k > m$ with $c_k \neq u_k$. This way we get 2 in the diagonal position of row k followed by $n - k$ 0's. Row $m < n$ is the only remaining reason for the matrix not to be lower triangular. To fix this, we use row n which is either all 1's or consists of $m + 1$ 0's followed by $n - m$ 2's. Adding half or one quarter of row n to row m , we get the matrix in lower triangular form. The row operations do not affect the determinant, which is now the product of the diagonal elements, which are all 1, $\frac{3}{2}$, or 2. This implies that the determinant is positive and therefore has the same sign as for the seed configuration, as claimed.

In the last step of the proof, we consider other choices for the centers of the C_k , reducing them to the above configuration, which we already know has positive orientation. Fix the set of indices k with $c_k \neq u_k$ and let m be the smallest such index, as before. We have c_{mj} equal to $\frac{1}{2}$ or $\frac{3}{2}$ for $j \leq m$ and equal to $-\frac{1}{2}$ or $\frac{1}{2}$ for $m < j$. Fixing c_m leaves only one choice for each $c_k \neq u_k$, else c_k and c_m would contradict (9). In the case we already studied, we had $c_{mj} \neq \frac{1}{2}$ for all j . The remaining cases use $\frac{1}{2}$ at least once as a coordinate. We claim that doing so does not change the determinant. We prove this by induction over the number of $\frac{1}{2}$'s in the coordinate vector of c_m . Each step decreases this number while preserving the set of rows for which

$c_k \neq u_k$. Let j be such that $c_{mj} = \frac{1}{2}$. Changing this coordinate to $-\frac{1}{2}$ or $\frac{3}{2}$, whichever is possible considering the value of u_{mj} , decreases the number of $\frac{1}{2}$'s, so it suffices to show that making that change does not affect the determinant. Indeed, the matrix before differs from the matrix after the change only in the j th column. Under the current assumptions, we have $c_{j-1} = u_{j-1}$, else c_{j-1} would be a vertex of U_j and C_{j-1} would contain U_j . Symmetrically, we get $c_j = u_j$. It follows that subtracting row $j-1$ from row j leaves only one non-zero element in row j , namely the 1 in column j . Using this row, we can now transform one matrix into the other by row operations, implying that the determinant does not change. Hence, the orientation of the n -simplex is the same as that of its seed configuration in all cases. $\square$

It is convenient to order the vertices of the simplices such that all n -simplices in $\mathcal{D}^n$ have positive orientation. Two neighboring n -simplices then induce opposite orientations on the shared $(n-1)$ -simplex.

Geometric Realization We are now ready to prove that dual complexes of balanced hierarchical subdivisions are simplicial complexes. To cope with the infinite size, we again consider finite subsets.

Geometric Realization Theorem *Let $\mathcal{C}$ be a balanced hierarchical cubical subdivision of $\mathbb{R}^n$. Then the dual complex $\mathcal{K}(\mathcal{C}(B))$ is geometrically realized in $\mathbb{R}^n$, for each cell $B \in \mathcal{C}$.*

Proof Let $\mathcal{R}_{\max}$ be the largest refinement of $\mathcal{C}$ with $\mathcal{C}(B) \subseteq \mathcal{R}_{\max}$. We can construct $\mathcal{R}_{\max}$ by repeatedly adding cubical cells on the outside of B , choosing the smallest size possible without violating balance. The layers of cells around B get smaller toward the outside until they shrink to unit size. Leaving two full layers of unit n -cubes, we remove all cubes outside those layers and consider $\mathcal{R}_{\max}$ in this extended box only. The two outer layers of unit cubes are useful because we understand how unit n -cubes are connected to each other in the dual complex. In particular, the full subcomplex defined by the subset of unit n -cubes in the two layers is geometrically realized in $\mathbb{R}^n$. Indeed, this is a subcomplex of $\mathcal{D}^n$, which we analyzed in Sect. 3.

For the final step of the argument, we compactify $\mathbb{R}^n$ to the n -dimensional sphere, $\mathbb{S}^n$, by adding a point at infinity. Similarly, we construct $\mathcal{K}^n$ from $\mathcal{K}(\mathcal{R}_{\max})$ by adding a new vertex at infinity and connecting it to all simplices triangulating the outer boundary. By the Nerve Theorem applied to the fractally distorted image, the thus modified dual complex of $\mathcal{R}_{\max}$ triangulates $\mathbb{S}^n$. It follows that the drawing of $\mathcal{K}(\mathcal{R}_{\max}) \subseteq \mathcal{K}^n$ in $\mathbb{R}^n$ defines a continuous mapping $g: \mathbb{S}^n \rightarrow \mathbb{S}^n$. We use the fact that the degree of g at a point x not in the image of any $(n-1)$ -simplex is the number of n -simplices that contain $g^{-1}(x)$, counting an n -simplex positive or negative depending on the orientation of its image under g ; see [1, p. 474] but also [13]. Since all cells in the last two layers are unit n -cubes, the n -simplices they define all have positive orientation. Hence, the degree is 1 if x lies inside the layer of n -simplices formed by the two layers of unit n -cubes. However, the degree of a mapping between manifolds without boundary is a global property and does not depend on the location of x ; see e.g. [1, p. 490]. Hence, it is 1 for any x . By the Orientation Lemma, the image under g of every n -simplex has positive orientation. Hence, the degree can only

be 1 if x lies in the interior of exactly one n -simplex. This prohibits improper intersections between simplices in $\mathcal{K}(\mathcal{R}_{\max})$. Since $\mathcal{K}(\mathcal{C}(B)) \subseteq \mathcal{K}(\mathcal{R}_{\max})$, this implies the claim. $\square$

6 Discussion

The main new concept in this paper is the dual complex of a cubical subdivision of $\mathbb{R}^n$. Important examples of the latter are quad-tree subdivisions of $\mathbb{R}^2$ and oct-tree subdivisions of $\mathbb{R}^3$. We count the number of simplices and prove that dual complexes of balanced hierarchical cubical subdivisions are geometrically realized in $\mathbb{R}^n$. We predict applications of these results in the analysis of four- and higher-dimensional images, and in particular in the computation of their persistent homology.

The detailed analysis of cubical subdivisions raises a number of technical questions. For example, the Geometric Realization Theorem applies only to balanced hierarchical cubical subdivisions. We know it does not necessarily hold for unbalanced such subdivisions of $\mathbb{R}^n$, for $n \geq 3$. How about balanced cubical subdivisions that are not hierarchical?

Acknowledgements This research is partially supported by the Defense Advanced Research Projects Agency (DARPA) under grants HR0011-05-1-0057 and HR0011-09-0065 as well as the National Science Foundation (NSF) under grant DBI-0820624.

References

1. Alexandroff, P., Hopf, H.: *Topologie I*. Springer, Berlin (1935)
2. Bendich, P., Edelsbrunner, H., Kerber, M.: Computing robustness and persistence for images. *IEEE Trans. Vis. Comput. Graph.* **16**, 1251–1260 (2010)
3. Bern, M., Eppstein, D., Gilbert, J.: Provably good mesh generation. *J. Comput. Syst. Sci.* **48**, 384–409 (1994)
4. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.: *Computational Geometry. Algorithms and Applications*, 3rd edn. Springer, Berlin (2008)
5. Dlotko, P., Kaczyński, T., Mrozek, M., Wanner, T.: Coreduction homology algorithm for regular CW-complexes. *Discrete Comput. Geom.* **46**, 361–388 (2011)
6. Edelsbrunner, H., Harer, J.L.: *Computational Topology. An Introduction*. Am. Math. Soc., Providence (2010)
7. Freudenthal, H.: Simplicialzerlegung von beschränkter Flachheit. *Ann. Math.* **43**, 580–582 (1942)
8. Fuchs, R., Hauser, H.: Visualization of multi-variate scientific data. *Comput. Graph. Forum* **28**, 1670–1690 (2009)
9. Johnson, C.: Top scientific visualization research problems. *IEEE Comput. Graph. Appl.* **24**, 13–17 (2004)
10. Kaczyński, T., Mischaikow, K., Mrozek, M.: *Computational Homology*. Springer, New York (2004)
11. Kuhn, H.W.: Some combinatorial lemmas in topology. *IBM J. Res. Dev.* **45**, 518–524 (1960)
12. Moore, D.: The cost of balancing generalized quadrees. In: *Proc. 3rd ACM Symp. Solid Modeling Appl.*, pp. 305–312 (1995)
13. Munkres, J.R.: *Elements of Algebraic Topology*. Perseus, Cambridge (1984)
14. Samet, H.: *The Design and Analysis of Spatial Data Structures*. Addison-Wesley, Reading (1990)
15. Sonka, M., Hlavac, V., Boyle, R.: *Image Processing, Analysis and Machine Vision*, 2nd edn. PWS, Pacific Grove (1999)

16. Stock, S.: High energy X-ray scattering quantification of bone strains. Talk at the First Workshop on “Bone Tissue: Hierarchical Simulations for Clinical Applications”, Los Angeles (2010)
17. Wagner, H., Chen, C., Vuçini, E.: Efficient computation of persistent homology for cubical data. In: Proc. 4th Workshop on Topology-based Methods in Data Analysis and Visualization (2011)
18. Weiss, K., De Floriani, L.: Diamond hierarchies of arbitrary dimension. *Comput. Graph. Forum* **28**, 1289–1300 (2009)