Wavelets on Irregular Point Sets

BY INGRID DAUBECHIES[†], IGOR GUSKOV[†],
PETER SCHRÖDER[‡], AND WIM SWELDENS[*]

[†] *Program for Applied and Computational Mathematics, Princeton University.*
[‡] *Department of Computer Science, California Institute of Technology.*
[*] *Lucent Technologies, Bell Laboratories.*

`ingrid@math.princeton.edu, ivguskov@math.princeton.edu, ps@cs.caltech.edu, wim@bell-labs.com`

In this article we review techniques for building and analyzing wavelets on irregular point sets in one and two dimensions. We discuss current results both on the practical and theoretical side. In particular we focus on subdivision schemes and commutation rules. Several examples are included.

1. Introduction

Wavelets are a versatile tool for representing general functions and data sets, and they enjoy widespread use in areas as diverse as signal processing, image compression, finite element methods, and statistical analysis (among many others). In essence we may think of wavelets as *building blocks* with which to represent data and functions. The particular appeal of wavelets derives from their representational and computational *efficiency*: most data sets exhibit correlation both in time (space) and frequency, as well as other types of structure. These can be modeled with high accurary through sparse combinations of wavelets. Wavelet representations can also be computed *fast*, because they can be built using multiresolution analysis and subdivision.

Traditionally, wavelet functions $\psi_{j,m}$ are defined as translates and dilates of one particular function, the *mother* wavelet ψ. We refer to these as *first generation* wavelets. This paper is concerned with a more general setting in which wavelets need not—and in fact, cannot—be translates and dilates of one or a few templates. Generalizations of this type were called *second generation* wavelets in [39]; they make it possible to reap the benefit of wavelet algorithms in settings with irregularly spaced samples, or on 2-manifolds which cannot be globally parameterized to the plane. In generalizing wavelet analysis to these more general settings one would like to preserve many of the properties enjoyed by first generation wavelets. In particular, they should still be associated with fast algorithms and have appropriate smoothness and localization properties. In addition, they should be able to characterize various functional spaces of interest. In this paper we shall be mostly concerned with fast algorithms, localization, and smoothness; we will not address function space characterizations. Note though that the smoothness of the wavelets is related to their ability to form unconditional bases for certain function spaces [8][17].

The key to generalizing wavelet constructions to these non-traditional settings is the use of generalized *subdivision schemes*. The first generation setting is already connected with subdivision schemes, but they become even more important in the construction of second generation wavelets. Subdivision schemes provide fast algorithms, create a natural multiresolution structure, and yield the underlying scaling functions and wavelets we seek.

Subdivision is a technique originally intended for building smooth functions starting from a coarse description. In this setting there is no need for irregular grids, as one is free to choose the finer grid to be regular. However, we intend to use subdivision as part of an entire multiresolution analysis which starts from a finest, irregular grid. This finest grid is gradually "coarsified"; subdivision then gives an approximation of the original data by extrapolating the reduced data on the coarser grid back to the original finest grid. In such a setting the geometry of the grids is fixed by the finest irregular grid and the coarsification procedure; thus subdivision on irregular grids is called for.
Note: Another approach would be to resample the original finest level data on a regular grid and use first generation wavelets. Resampling however can be costly, introduce artefacts, and is generally impossible in the surface setting. Therefore we choose to work on the original grid.

(*a*) *1D Subdivision*

The main idea behind subdivision is the iteration of upsampling and local averaging to build functions and intricate geometrical shapes. Originally such schemes were studied in computer aided geometric design in the context of corner cutting [14][7] and the construction of piecewise polynomial curves, e.g., the de Casteljau algorithm for Bernstein-Bézier curves [13], and algorithms for the iterative generation of splines [31][1]. Later subdivision was studied independently of spline functions [21][19][15][4][5][6] and the connection to wavelets was made [33][9].

For example, Figure 1 demonstrates the application of the four point scheme. New points are defined as local averages of two old points to the left and two old points to the right with weights $(-1, 9, 9, -1)/16$.

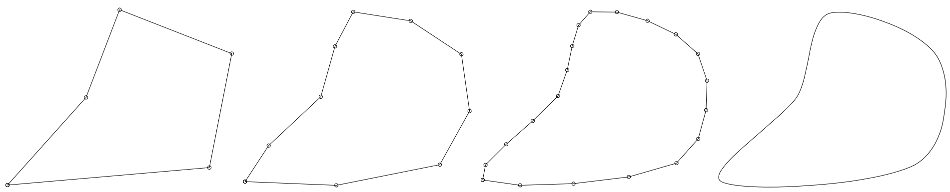

Figure 1. Subdivision is used to generate a smooth curve starting from a coarse description.

In the case of spline functions, smoothness follows from simple algebraic conditions on the polynomial segments at the knots. However, in the general setting convergence and smoothness of the limit function are harder to prove. Various approaches have been explored to find the Hölder exponent of the limit function, or to determine its Sobolev class. Early references in this context are [19][21][15][34][12][23][4] [36][41][24]. These studies and their results all rely on regular, i.e., equi-spaced grids. The analysis uses tools such as the Fourier transform, spectral analysis, and the commutation formula.

In this paper we focus on irregular point sets. To describe the settings we are interested in, we distinguish three types of refinement grids: regular, semi-regular, and irregular; see Figure 2. A *regular* grid has equidistant points on each level and each time a new point is inserted exactly between two old points. For example, the curve shown in Figure 1 is parameterized over a regular grid. A *semi-regular* grid (middle in Figure 2) starts with an irregular coarse grid and adds new points at parameter locations midway between successive old points. Thus the finer grids are locally regular except around the original coarsest level points. In *irregular* grids (right) parameter locations of new points need not be midway between successive old points. Note that regular grids are translation and dilation invariant, semi-regular grids are locally dilation invariant around coarsest level vertices, and irregular grids generally possess no invariance property.

Similarly the weights used in subdivision schemes come in three categories: uniform, semi-uniform, and non-uniform. *Uniform* schemes like the four point scheme of Figure 1 correspond to first generation wavelets; and use the same subdivision weights within a level and across all levels; they are typically used on regular grids or grids which can be smoothly remapped to a regular grid. *Semi-uniform* schemes are used on semi-regular grids; they vary the weights within each level (special weights are used in the neighborhood of the coarsest level points), but the same weights are used across levels. Such schemes are sometimes referred to as *stationary*. Wavelets and subdivision schemes on an interval also fall in this category. *Non-uniform* schemes use varying weights within *and* across levels and correspond to the second generation setting.

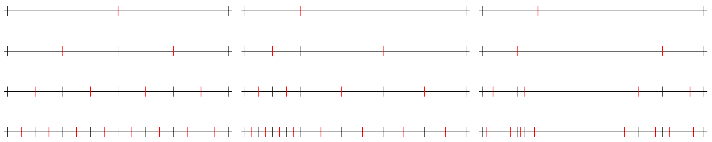

Figure 2. Regular, semi-regular, and irregular grid hierarchies in 1D.

Almost all work on smoothness for non regular grids concerns the semi-regular grids with semi-uniform subdivision schemes. Because translation-invariance is lost, Fourier transform based arguments can no longer be used. However, since the same weights are used on successive levels, one has dilation invariance around coarsest level points and can reduce the smoothness analysis to the study of spectral properties of certain fixed matrices. In [42][43] Warren shows that the semi-uniform version of the four point scheme on a semi-regular grid yields a C^1 limit function.

In the irregular case the subdivision scheme must become non-uniform to account for the irregularity of the associated parameter locations. This is illustrated in Figure 3 which shows the limit functions of the uniform four point rule (left) and non-uniform four point rule [40] (right); both use the same irregular grid.

Figure 3. An example why non-uniform subdivision is needed. Left the limit function with uniform subdivision, right non-uniform. The same irregular grid is used in both figures.

The study of irregular subdivision is not only theoretically interesting, but also of great importance in practical applications. For example, in the semi-regular setting, one can use adapted weights to better control the shape of a curve [30] or surface [45]. More importantly, in many practical setups we start with samples associated to a very fine, but irregular grid. Now the main task for subdivision is not further refinement but rather aid in a multiresolution analysis on *coarser* grids. The wavelet and scaling functions from the coarsest level are generated with a subdivision scheme with new points which are no longer parametric midpoints, but are dictated by the finest level grid on which the data was originally sampled. Even though the actual number of levels is always finite for any concrete application of these methods, the asymptotic behavior of irregular subdivision is still relevant as the finest and coarsest level could be arbitrarily far apart.

In these settings smoothness results become much harder to obtain. Because the subdivision weights vary within a level, the Fourier transform can no longer be used, and because they vary across levels even spectral analysis cannot help. In this paper we discuss some tools that can be used to analyze smoothness; in particular we demonstrate that the commutation formula still holds and becomes a critical tool for smoothness analysis.

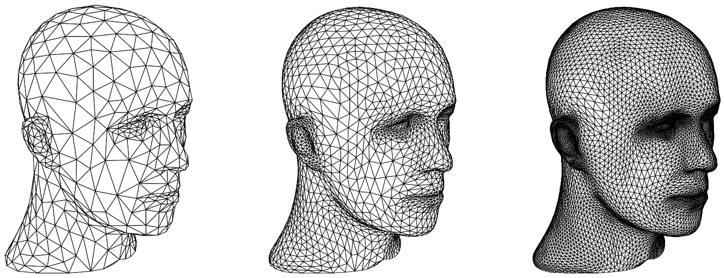

Figure 4. 2D Loop Subdivision is used to generate a smooth surfaces from a coarse description.

(b) 2D Subdivision

The 2D setting appears in the context of generating smooth surfaces, see Figure 4. Here regular grids are too restrictive. For example, tensor product settings are only applicable for surfaces homeomorphic to a plane, cylinder, or torus due to the Euler characteristic. Historically this challenge was addressed by generalizing traditional spline patch methods to the semi-regular setting [18] (bi-quadratic) [3] (bi-cubic) [32] (quartic box-spline). Similar to the 1D setting researchers also developed interpolating constructions [22][45][28]. All these settings (and others since; for an overview see [37]) proceed by applying quadrisection to an initial mesh consisting of either quadrilaterals or triangles and thus belong to the semi-regular setting. The weights used in the subdivision scheme are semi-uniform since they take into account the local neighborhood structure of a vertex, i.e., how many edge neighbors a given vertex has. As in the 1D semi-regular setting spectral analysis is the key to understanding the smoothness of these constructions. We refer to [35][42][44][38] for more details.

The irregular setting appears in 2D just as in the 1D case when some finest irregular level is presented on input and the main task is to build a multiresolution analysis on *coarser* levels. In this case however, we can no longer define downsampling as simply retaining every other sample. This brings us to the realm of mesh simplification; we postpone the discussion of mesh simplification and the construction of appropriate non-uniform subdivision operators to Section 3.

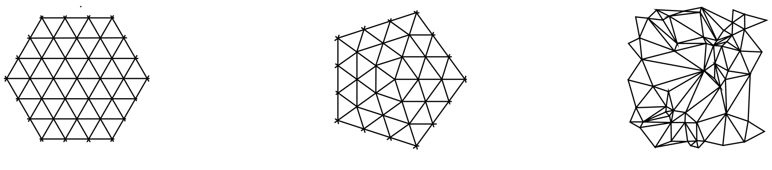

Figure 5. Sections of regular, semi-regular, and irregular triangle grids in 2D.

(*c*) *Overview*

This paper summarizes results obtained in [11][10][25][26]. We start with the onedimensional results of [11][10]. We show that even simple subdivision rules, such as cubic Lagrange interpolation, can lead to very intricate subdivision operators. To control these operators, we use commutation: because the subdivision scheme maps the space of cubic polynomial sequences to itself, we can define derived subdivision schemes for the divided difference sequences. These simpler schemes can then be used to prove growth bounds on divided differences of some order, corresponding to smoothness results for the limit function of the original scheme. The commutation formula enables us to control smoothness and is the key to the construction of wavelets associated with the subdivision scheme.

In [25] inspiration from the 1D analysis is used to tackle the much more complex 2D case. Again, differences and divided differences are introduced, which can be computed from level to level with their own derived subdivision scheme. Control on the growth of these divided differences then leads to smoothness results. In practice, finding the right Ansatz for irregular subdivision in the 2D setting is much harder than in the already difficult 1D case. Finally, we show how irregular subdivision schemes can be used in multiresolution pyramids for 2D meshes embedded in $\mathbf{R}^3$ and review several applications from [26]. The "wavelets" associated with these schemes are overcomplete and are related to frames rather than bases.

2. The one dimensional case

(*a*) *Multilevel grids*

Consider grids X_j, which are strictly increasing sequences of points $\{x_{j,k} \in \mathbf{R} \mid k \in \mathbf{Z}\}$, and which are consecutive binary refinements of the initial grid X_0, i.e., $X_j \subset X_{j+1}$ and $x_{j+1,2k} = x_{j,k}$ for all j and k. Thus in every refinement step we insert one odd indexed point $x_{j+1,2k+1}$ between each adjacent pair of "even" points $x_{j,k} = x_{j+1,2k}$ and $x_{j,k+1} = x_{j+1,2k+2}$, as in Figure 2. We define $d_{j,k} := x_{j,k+1} - x_{j,k}$. We shall also use the term *grid size* on level j, for the quantity $d_j := \sup_k d_{j,k}$. As $j \to \infty$ we want the grids to become dense, "with no holes left"; this translates to the requirement that the d_j be summable.

Notes:

(i) The above multilevel grids are called *two-nested*. One can also consider more general irregular grids such as q-nested grids where we insert $q-1$ new points in between old points or even non-nested but "threadable" grids. See [10] for more details on this.

(ii) In case the ratio between the lengths of any two neighboring intervals is globally bounded, we call the grid *homogeneous*. An example of an irregular two-nested grid that is not homogeneous is built by $x_{j+1,2k+1} = \beta x_{j,k} + (1-\beta)x_{j,k+1}$ where β is a fixed parameter satisfying $0 < \beta < 1$. This is an example of a *dyadically* balanced grid: the ratio between the lengths of two "sibling" intervals $d_{j,2l}$ and $d_{j,2l+1}$ is bounded. However the ratio between $d_{j,-1} = \beta^j$ and $d_{j,0} = (1-\beta)^j$ is unbounded.

(*b*) *Subdivision schemes*

Subdivision starts with a set of initial function values $f_0 = \{f_{0,k}\}$ which live on the coarsest grid X_0. The subdivision scheme S is a sequence of linear operators S_j, $j \geqslant 0$, which iteratively compute values $f_j = \{f_{j,k}\}$ on the finer grids via the formula

$f_{j+1} = S_j\, f_j$, or

$$f_{j+1,l} = \sum_k S_{j,l,k}\, f_{j,k}.$$

We consider only *local* schemes in the sense that the above summation has a globally

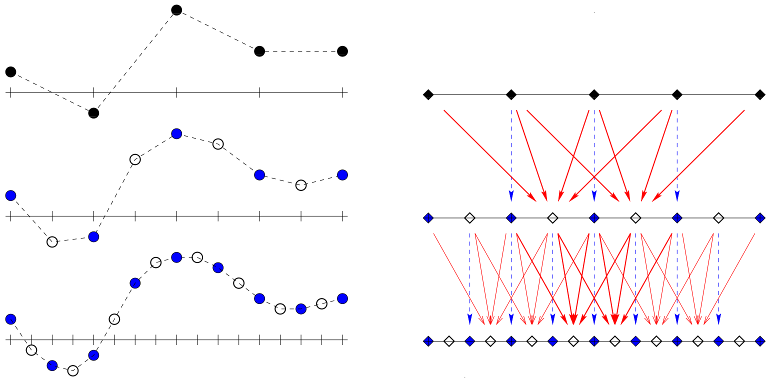

Figure 6. On the left an interpolating scheme is applied to function values; the new function values are shown as open circles. On the right arrows show the dependencies in the computation of those values; the vertical dashed arrows indicate that function values that were already assigned are kept unchanged in the subdivision, because this is an interpolating scheme.

bounded number of terms centered around $k = 2l$. Subdivision gives us values defined on the grid points $x_{j,k}$. By connecting these points we can define a piecewise linear function $f_j(x)$, see Figure 6. Our ambition is to synthesize a continuous limit function $\varphi(x)$ as the pointwise limit for $j \to \infty$ of $f_j(x)$. We are interested in the existence and smoothness of $\varphi(x)$.

The subdivision coefficients $S_{j,k,l}$ will depend on the application one has in mind. We pointed out in the introduction that one cannot simply stick with the coefficients from the regular case; typically the coefficients need to be spatially varying, and will be linked to the spatial variation of the grid.

One such subdivision scheme which allows for a spatial interpretation is Lagrangian interpolating subdivision [19][21][15][16]. Here the value $f_{j+1,2k+1}$ at a new point is found by defining a polynomial which interpolates the points $(x_{j,l}, f_{j,l})$ for l in the neighborhood of k, and evaluating this polynomial at $x_{j+1,2k+1}$; see Figure 7. In the reg-

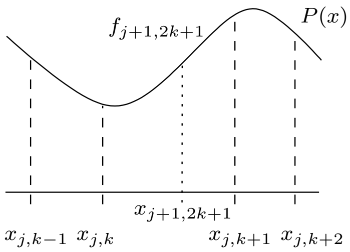

Figure 7. Cubic Lagrangian interpolation: The value $f_{j+1,2k+1}$ at the odd gridpoint $x_{j+1,2k+1}$ is obtained by evaluating a cubic polynomial $P(x)$ interpolating values at 4 neighboring even gridpoints $x_{j+1,2k-2} = x_{j,k-1}, \ldots, x_{j+1,2k+4} = x_{j,k+2}$.

ular cubic case, this corresponds to the standard four point scheme, with $S_{j,2k+1,k} =$

$S_{j,2k+1,k+1} = 9/16$ and $S_{j,2k+1,k-1} = S_{j,2k+1,k+2} = -1/16$. In the irregular setting the coefficients are a non-trivial quotient of cubic polynomials in the $x_{j,k}$, see [11].

Lagrangian subdivision is *interpolating* in the sense that in each subdivision step the values at the even grid points are kept, i.e. $f_{j+1,2k} = f_{j,k}$, and the limiting function thus interpolates the original data $\varphi(x_{0,k}) = f_{0,k}$. For non-interpolating or *approximating* schemes the $f_{j+1,2k}$ can differ from $f_{j,k}$ (see Figure 8).

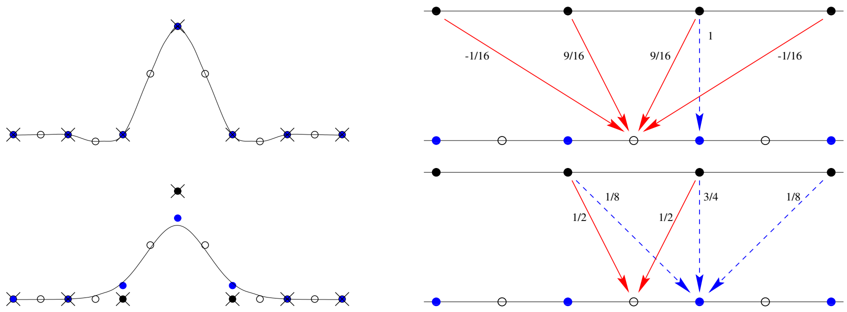

Figure 8. The top row shows the limit function and weights of the cubic Lagrange interpolation scheme, the bottom row illustrates the non-interpolating subdivision scheme producing cubic B-splines.

(*c*) *Smoothness results*

To derive smoothness estimates, we use Lemarié's *commutation formula* idea, generalized to the present irregular setting. (Note that this is similar to [36] [20] which study the regular case.) For the cubic Lagrange interpolation example, this amounts to introducing divided difference sequences

$$f^{[1]}_{j,k} := \frac{f_{j,k+1} - f_{j,k}}{x_{j,k+1} - x_{j,k}},$$

and observing that the $f^{[1]}_{j,k}$ are *also* related by local subdivision, i.e. there exists a local subdivision scheme with entries $S^{[1]}_{j,l,k}$ so that

$$f^{[1]}_{j+1,l} = \sum_k S^{[1]}_{j,l,k} f^{[1]}_{j,k}.$$

The existence of the $S^{[1]}_j$ follows from the fact that every $S^{[0]}_j := S_j$ maps a constant sequence to itself (see [11]). (It is clear that if $S^{[0]}$ does not leave constants invariant, then $S^{[1]}$ cannot exist, since it would need to map the zero sequence to a non-zero result.) Moreover, one can show that if the $f^{[1]}_j$ "converge" to a continuous function φ_1, then the $f^{[0]}_j := f_j$ converge as well, to a continuously differentiable function φ, and that $\varphi' = \varphi_1$. (For details, see [11] or [10].) This is the essence of the commutation idea:

$$\begin{array}{ccc} f_j & \xrightarrow{\text{divided difference}} & f^{[1]}_j \\ \Big\downarrow \text{limit} & & \Big\downarrow \text{limit} \\ \varphi & \xrightarrow{\text{differentiation}} & \varphi_1 = \varphi' \end{array}$$

It turns out that one can also consider higher order divided differences; for the cubic

Lagrange interpolation case, one can go up to fourth order differences because $S_j^{[0]}$ maps cubic polynomials sampled at $x_{j,k}$ to cubic polynomials sampled at $x_{j+1,k}$. These $f_{j,k}^{[4]}$ no longer converge, but we can control their growth, and this helps us prove that φ_1 is continuous, and φ continuously differentiable. In fact, detailed (and rather technical) estimates in [11] show that, for homogeneous grids,

$$\left| f_{j,k}^{[4]} \right| \leqslant C \frac{\lambda^j}{d_{j,k}^3},$$

where $\lambda < 1$ is determined by the bound on the ratio between neighboring interval lengths. Once such a bound is known, a general theorem (see Theorem 4 in [11]) can be used to prove that $\varphi \in C^{2-\epsilon}$. This result is optimal in the sense that even in the regular case, better smoothness cannot be obtained.

Notes:

(i) This result for cubic Lagrange interpolation on homogeneous grids can be extended to grids that are dyadically balanced only. The analysis becomes much more delicate.

(ii) A similar approach can be used for non-interpolating subdivision. In that case it turns out that one has to use appropriately defined divided differences, which are different from the "standard" definition. See [10] for a complete discussion of this situation.

(*d*) *Wavelets*

Wavelets at level j are typically used, in the regular case, as the building blocks to represent any function in the multiresolution analysis that "lives" in the $(j+1)$-st approximation space V_{j+1}, but not in the coarser resolution approximation space $V_j \subset V_{j+1}$. One can introduce similar wavelets in the present irregular setting. The scaling functions $\varphi_{j,k}$ are the limit functions obtained from starting the subdivision scheme at level j, from the "initial" data $f_{j,l} = \delta_{l,k}$, and refining from there on. Under appropriate assumptions on the subdivision operators S_j, the $\varphi_{j,k}$ are independent; V_j is the function space spanned by them. Clearly, $V_j \subset V_{j+1}$. As in the regular case, there are many different reasonable choices for complement spaces W_j (which will be spanned by the wavelets at level j) that satisfy $V_{j+1} = V_j \oplus W_j$.

When the scaling functions are interpolating as in the Lagrangian case, i.e., $\varphi_{j,k}(x_{j,k'}) = \delta_{k,k'}$, then a simple choice for a wavelet is given by $\psi_{j,m} = \varphi_{j+1,2m+1}$, i.e., the wavelet is simply a finer scale scaling function at an odd location. This is sometime called an *interpolating* wavelet. This is in general not a very good wavelet as it does not have any vanishing moments. It can be turned into a wavelet with vanishing moments using the lifting scheme [40].

Another way to select a complement space W_j, is to use commutation between two biorthogonal multiresolution hierarchies, V_j and $\tilde{V}_j$. If both are associated to local subdivision schemes, then the biorthogonality of the $\varphi_{j,k}$ and $\tilde{\varphi}_{j,l}$ imposes consistency requirements on the S_j and $\tilde{S}_j$. Commutation can be used, as in the regular case, to pass from one dual pair of multiresolution analyses to another, by operations related to differentiating and integrating, respectively. For instance, the above choice of an interpolating wavelet corresponds formally to letting the dual scaling function be a Dirac. Applying the commutation rule each time reduces the order of the scaling functions, but increases the order of the dual scaling function. In particular the Dirac will become a box and later on a general B-spline. It turns out that there is a natural definition of wavelets $\psi_{j,k}$ and

$\tilde{\psi}_{j,k}$ corresponding to the dual multiresolution structures. It is shown in [10] that as in the regular case the new wavelet after commutation is the derivative of the old wavelet and the new dual wavelet is the integral of the old dual wavelet. By repeatedly applying commutation starting from the Lagrangian setting one can thus build the entire family of biorthogonal, compactly supported, irregular B-spline wavelets and their duals [10].

3. The two dimensional case

We mentioned in the introduction that the importance of the irregular setting arises from the practical need to coarsify in settings in which the initial input is given as a function over a fine triangulation of the plane (functional setting) or as a triangulation of a 2-manifold (surface setting). In the 1D setting the downsampling operation to create a coarser level is straightforward as we can simply "skip" every other sample. In the irregular 2D setting downsampling is much less straightforward. Before delving into the details of irregular 2D subdivision we first discuss a number of approaches which can be employed to define irregular downsampling in the surface setting. This problem has received a lot of attention in computer graphics where it is generally referred to as polygonal simplification.

(*a*) *Polygonal Simplification*

In polygonal mesh simplification the goal is to simplify a given (triangulated) mesh $\mathcal{M}^L = (\mathcal{P}^L, \mathcal{K}^L)$ into successively coarser, homeomorphic meshes $(\mathcal{P}^l, \mathcal{K}^l)$ with $0 \leqslant l < L$, where $(\mathcal{P}^0, \mathcal{K}^0)$ is the coarsest or base mesh. Here $\mathcal{P}^l$ is a set of l point positions, while $\mathcal{K}^l$ encodes the topological structure of the mesh and consists of triples $\{i, j, k\}$ (triangles), pairs $\{i, j\}$ (edges), and singletons $\{i\}$ (vertices). The goal now is to allow certain topological operations on $\mathcal{K}^l$ which preserve the manifold property and genus of the mesh. These changes go hand in hand with geometric changes which are typically subject to an approximation quality criterion.

Several approaches for such mesh simplification have been proposed (the interested reader is referred to the excellent survey [27] for more details). The most popular methods are so called "progressive meshes" (PM). In a PM contruction a sequence of edge collapses is prioritized based on the error it introduces. An edge collapse brings the endpoints of the chosen edge into coincidence, in the process removing two triangles, three edges, and one vertex (in the case of interior edges). The point location of the merged vertex can be chosen so as to minimize some error criterion with respect to the original mesh. The error can be measured in various norms such as L^∞ (symmetric Haussdorff distance), L^2, and Sobolev norms.

For our purposes we are using a PM construction based on half edge collapses, i.e., the point position for the collapsed edge is one of its end points. This results in a mesh hierarchy which is interpolating in the sense that the point position sets $\mathcal{P}^l$ are nested. There are several possible ways to define levels of a hierarchy. The most flexible one treats a single half edge collapse operation as defining a level. In contrast to the usual wavelet setting this results in a linear, rather than logarithmic, number of levels.

Before going to the surface case we first consider the functional setting and then treat the surface setting as three instances of a functional setting.

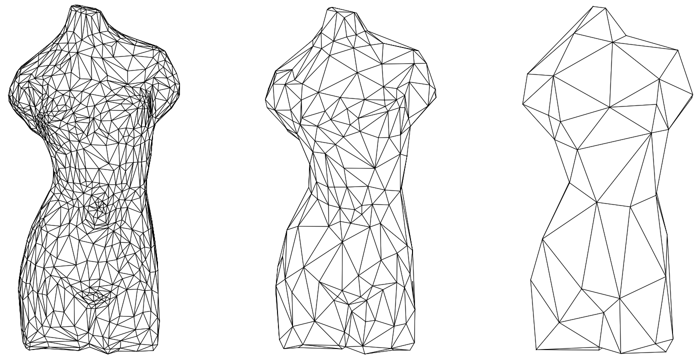

Figure 9. If an irregular, finely detailed mesh is given, the first task in building a multiresolution analysis is coarsification.

(*b*) *Functional Setting: Multivariate commutation formula*

Just as in the one dimensional case, irregular multivariate subdivision schemes act on sequences whose elements are associated with irregular parameter locations. We introduce levels numbered $0, 1, \ldots$ with level 0 corresponding to the coarsest scale. Within each level n the collection of all parameter locations constitute an irregular grid $\mathcal{X}_n$.

We can now introduce a subdivision scheme S as a sequence of linear operators S_n, $n \geqslant 0$ which iteratively compute sequences f_n defined on $\mathcal{X}_n$, starting from some coarsest level data f_0, via

$$f_{n+1} = S_n f_n.$$

In the one-dimensional setting we analyzed the regularity of the functions produced by subdivision through the behaviour of properly defined divided differences. We proceed similarly for the irregular two-dimensional setting. Let $\mathcal{D}_n^{[p]}$ denote the operator which maps the data sequence f_n into the corresponding sequence $f_n^{[p]}$ of divided differences of order p , that is, $f_n^{[p]} = \mathcal{D}_n^{[p]} f_n$. We say that there exists a derived subdivision scheme $S^{[p]}$ satisfying the commutation formula if the sequences $f_n^{[p]}$ are related via the relation $f_{n+1}^{[p]} = S_n^{[p]} f_n^{[p]}$ where the $S_n^{[p]}$ constitute a local, bounded subdivision scheme. Thus we can write $\mathcal{D}_{n+1}^{[p]} S_n = S_n^{[p]} \mathcal{D}_n^{[p]}$. We then prove that the bounds on the growth of sequences $f_n^{[p]}$ can be translated into the smoothness estimates for the functions produced by the original subdivision scheme S.

In order to extend this construction to the multivariate case we need to define the multivariate divided differences in such a way that the algebra of the commutation formula works. This is done in [25] for a class of polynomial reproducing subdivision schemes. It is also shown there, that for multilevel grids satisfying some natural conditions, the bounds on the growth of these divided differences can be used to analyze the regularity of functions produced by subdivision.

(*c*) *Constructing a subdivision scheme*

In this section we provide a particular example of a subdivision scheme which in the functional setting produces visually smooth functions on irregular triangulations. Our subdivision algorithm relies on minimizing divided differences. Consider a triangle $\{i, j, k\}$ in the parameter plane with corners (x_i, y_i), (x_j, y_j), and (x_k, y_k) and function values f_i, f_j, and f_k. These three function values define a plane. The gradient to this

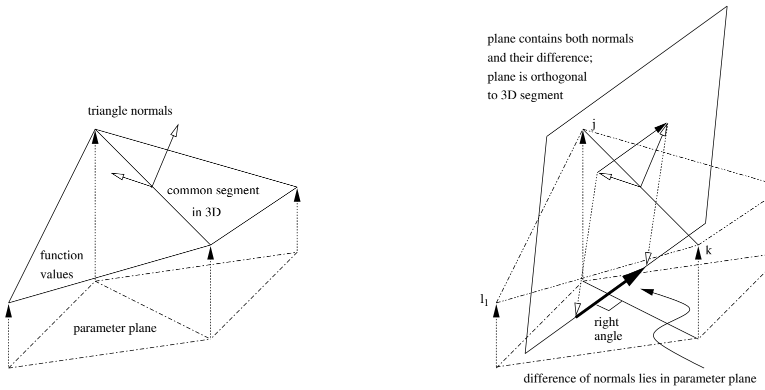

Figure 10. Second differences are associated with an edge. Since they are the difference of two adjacent triangle normals (first dividided differences) one can see that the second differences are orthogonal to the common edge in the parameter plane.

plane can be seen as a first order divided difference corresponding to this triangle. The gradient is zero only if the plane is horizontal ($f_i = f_j = f_k$).

Next we define the second order differences. They are computed as the difference between two normals on neighboring triangles and can be thought of as associated with the common edge (see Figure 10, left). It is easy to see that the difference between gradients of two adjacent triangles is orthogonal to their common edge (see Figure 10, right). Thus the component $D_e^2 f$ normal to the edge e can be used for the second order difference. It depends linearly on the four function values of these two triangles. The coefficients can be found in [25] or [26]. The second order difference operator is zero only if the two triangles lie in the same plane, and one can see that its behavior is closely related to the dihedral angle.

The central ingredient in the design of our subdivision scheme is the use of a non-uniform relaxation operator which inserts new values in such a manner that second order differences are minimized. Define a quadratic energy, which is an instance of a discrete fairing functional [29]

$$Rf_i = \text{arg}\ \min \mathbf{E}(f_i) = \sum_{e \in \mathcal{K}} (D_e^2 f)^2.$$

Setting $\partial \mathbf{E} / \partial f_i = 0$ yields

$$Rf_i = \sum_{j \in \mathcal{V}_2(i)} w_{i,j} f_j \quad \text{with } V_2(i) = \tag{3.1}$$

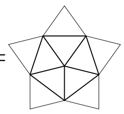

Note that if f is a linear function, i.e., all triangles lie in one plane, the fairing functional $\mathbf{E}$ is zero. Consequently linear functions are invariant under R. In particular R preserves constants from which we deduce that the $w_{i,j}$ sum to one.

Subdivision is computed one level at a time starting from level n_0 in the PM. Reversing the PM construction back to the finest level adds one vertex (x_n, y_n, f_n) per level; the non-uniform subdivision is computed one vertex at a time. The position of each new vertex n is computed according to (3.1) using areas and lengths of the original finest level mesh. Next the immediate neighbors of n are relaxed using (3.1) as well. The ambition of our strategy of minimizing $D_e^2 f$ is to obtain C^1 smoothness. However, there is currently no Ansatz on the bounds of the divided differences to prove regularity of the limit function. Figure 11 shows an irregular grid of 20493 triangles (left), simplified down to 86 triangles (middle). Now associate the value $f = 1$ with the center vertex and 0 with all others. The figure on the right is the result of running the subdivision scheme back to the finest level. Even though the grid is irregular the resulting function appears smooth.

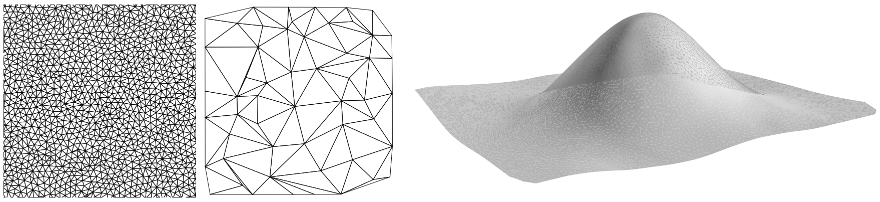

Figure 11. From left to right: a portion of the fine mesh; the coarse mesh; function produced by the non-interpolating scheme.

(d) Functions on surfaces

In order to build a multiresolution structure on meshes we first need to introduce the relaxation operator acting on functions defined over triangulated surfaces in 3D. We shall follow the strategy of the planar case and introduce second differences for such functions. For this we need to specify some locally consistent parameterization over the support of the difference operator. Consider a triangular mesh $\mathcal{P}$ in $\mathbf{R}^3$, and let $f : \mathcal{P} \to \mathbf{R}$. We would like to define the second difference operator $D_e^2 f$ for an edge e from the triangulation $\mathcal{P}$. For this we only need a consistent parametrization (i.e. flattening) for two neighboring triangles at a time. Let the edge $e = \{i, j\}$ be adjacent to two triangles $\{i, j, k\}$ and $\{j, i, l\}$. We use the "hinge map" to build a pair of adjacent triangles in the plane. These two triangles in the parameter plane have the same angles and edge lengths as the two triangles in $\mathbf{R}^3$. We then define $D_e^2 f$ as described in the previous section. Using these second differences, it is easy to extend the definition of the relaxation operator and the corresponding subdivision scheme to work with functions defined over triangulated surfaces.

(e) Burt-Adelson Pyramid

For meshes we found it more useful to generalize an oversampled Burt-Adelson type pyramid [2] than a critically sampled wavelet pyramid. Let $(\mathcal{P}^n)$ be some fixed PM hierarchy of triangulated surfaces. We start from the function $f_N : \mathcal{P} = \mathcal{P}^N \to \mathbf{R}$ defined on the finest level and compute a sequence of functions (f_n) $(n_0 \leqslant n \leqslant N)$ as well as oversampled differences $d_i^{(n)}$ between levels.

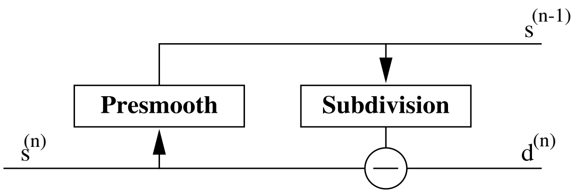

Figure 12. Burt-Adelson style pyramid scheme.

Like subdivision, the Burt-Adelson pyramid is computed vertex by vertex. Thus the four critical components of a BA pyramid: presmoothing, downsampling, subdivision, and detail computation are done for one vertex n at a time (see Figure 12). The presmoothing comes down to applying the relaxation operator to the neighbors of n. Downsampling simply removes the vertex n through a half edge collapse. We perform subdivision as described above and compute details $d^{(n)}$ for all neighbors of n.

To see the potential of a mesh pyramid in applications it is important to understand that the details $d^{(n)}$ can be seen as an approximate frequency spectrum of the mesh. The details $d^{(n)}$ with large n come from edge collapses on the finer levels and thus correspond to small scales and high frequencies, while the details $d^{(n)}$ with small n come from edge collapses on the coarser levels and thus correspond to large scales and low frequencies. Hence the sequence of $d^{(n)}$ for running n can be seen as an approximate frequency spectrum. Moreover, while the superscript n of an individual detail vector $d_i^{(n)}$ corresponds to its level/frequency, the subscript i corresponds to its location. Thus we actually have a space-frequency decomposition.

It is theoretically possible to build a critically sampled wavelet based on the Lifting scheme [39]. The idea is to use an interpolating subdivision scheme which only affects the new vertex and omits the relaxation of the even vertices. Consequently only one detail per vertex is computed and the sampling is always critical. However, at this point it is not clear how to design updates that make the transform numerically stable. Additionally, interpolating subdivision schemes do not yield very smooth meshes and have unwanted undulations. Therefore critically sampled wavelet transforms have had limited use in graphics applications.

4. Applications

In the *surface setting* we deal with a triangulated mesh $\mathcal{P}$ of arbitrary topology and connectivity embedded in 3D with vertices $p_i = (x_i, y_i, z_i)$. It is important to separate the two capacities the mesh $\mathcal{P}$ fulfills in our analysis. First, the original mesh and its PM representation serve as the source of local parameterization and connectivity information which determines the coefficients of our adaptive relaxation operator.

Second, if our purpose is to process the geometry of the mesh, it is crucial to treat all three coordinates x, y, and z as *dependent* variables. In fact, we consider the coordinates of the mesh to be real functions on the current PM vertex set. Initially, before any changes in geometry take place, these functions can be viewed as identities. When the wanted processing operations such as filtering or editing are applied to the data these functions become more meaningful.

As an example of possible application of our scheme we present various manipulations of the scanned Venus' head model. The original mesh has 50,000 vertices. After building a PM hierarchy, we use our BA pyramid scheme to build a multiresolution representation. We can use different manipulations of the detail coefficients in order to achieve various signal processing tasks. Specifically, if all the detail coefficients finer than some level are put to zero, we achieve a smoothing effect (in Figure 13(b) all the details on the levels above 1000 were set to zero). The stopband filter effect is achieved by setting to zero some range of coefficients (in Figure 13(c) all the details between the levels 1000 and 15000 were set to zero). One can also enhance certain frequencies (in Figure 13(d) all the details between the levels 1000 and 15000 were multiplied by two).

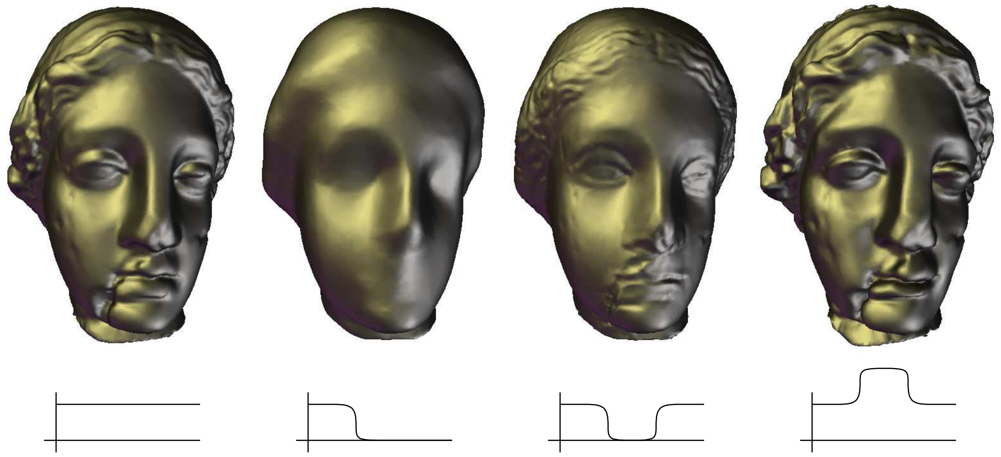

Figure 13. Smoothing and filtering of the venus head. From left to right: (a) original; (b) low pass filter; (c) stopband filter; (d) enhancement filter.

5. Conclusion

One of the current frontiers in wavelet research and applications is the generalization of multiresolution methods from the regular to the semi-regular and, more recently, irregular setting. We have given a brief review of these developments, starting with the one dimensional setting and moving on to the two dimensional functional and manifold settings. While there exists an extensive set of tools for the analysis of wavelet constructions in the regular setting, such tools have only recently begun to emerge for the irregular setting. One such tool is the generalization of commutation from the regular to the irregular setting. We have applied these ideas by proposing new irregular subdivision schemes in the manifold setting which are explicitly designed to minimize certain differences. Little is as yet known about the analytic smoothness properties of the resulting constructions, but numerical evidence suggests that they are quite useful for practical applications.

Acknowledgements

The work reviewed in this paper was supported in part by grants from NSF (ACI-9624957, ACI-9721349, DMS-9874082, DMS-9872890), DOE (W-7405-ENG-48), ONR (N00014-96-1-0367 P00004) and AFOSR (F49620-98-1-0044). Igor Guskov was partially supported by a Harold W. Dodds Fellowship and a Summer Internship at Bell Laboratories, Lucent Technologies.

References

[1] C. De Boor. A practical guide to splines. Number 27 in Applied Mathematical Sciences. Springer, New York, 1978.

[2] P. J. Burt and E. H. Adelson. Laplacian pyramid as a compact image code. IEEE Trans. Commun., 31(4):532–540, 1983.

[3] E. Catmull and J. Clark. Recursively generated B-spline surfaces on arbitrary topological meshes. Computer Aided Design, 10(6):350–355, 1978.

[4] A. S. Cavaretta, W. Dahmen, and C. A. Micchelli. Stationary subdivision. Memoirs Amer. Math. Soc., 93(453), 1991.

[5] A. S. Cavaretta and C. A. Micchelli. Computing surfaces invariant under subdivision. Computer Aided Geometric Design, 4(4):321–328, 1987.

[6] A. S. Cavaretta and C. A. Micchelli. The desing of curves and surfaces by subdivision algorithms. In T. Lyche and L. L. Schumaker, editors, Mathematical Aspects of Computer Aided Geometric Design. Academic Press, Tampa, 1989.

[7] G. Chaikin. An algorithm for high speed curve generation. Comp. Graphics Image Process., 3:346–349, 1974.

[8] W. Dahmen. Stability of multiscale transformations. J. Fourier Anal. Appl., 2(4):341–361, 1996.

[9] I. Daubechies. Orthonormal bases of compactly supported wavelets. Comm. Pure Appl. Math., 41:909–996, 1988.

[10] I. Daubechies, I. Guskov, and W. Sweldens. Commutation for irregular subdivision. Technical Report, Bell Laboratories, Lucent Technologies, 1998.

[11] I. Daubechies, I. Guskov, and W. Sweldens. Regularity of irregular subdivision. Constr. Approx., 15:381-426, 1999.

[12] I. Daubechies and J. C. Lagarias. Two-scale difference equations I. Existence and global regularity of solutions. SIAM J. Math. Anal., 22(5):1388–1410, 1991.

[13] F. de Casteljau. Outillages méthodes calcul. André Citroën Automobiles SA, Paris, 1959.

[14] G. de Rham. Sur une courbe plane. J. Math. Pures Appl., 39:25–42, 1956.

[15] G. Deslauriers and S. Dubuc. Interpolation dyadique. In Fractals, dimensions non entières et applications, pages 44–55. Masson, Paris, 1987.

[16] G. Deslauriers and S. Dubuc. Symmetric iterative interpolation processes. Constr. Approx., 5(1):49–68, 1989.

[17] D. L. Donoho. Interpolating wavelet transforms. Preprint, Department of Statistics, Stanford University, 1992.

[18] D. Doo and M. Sabin. Analysis of the behaviour of recursive division surfaces near extraordinary points. Computer Aided Design, 10(6):356–360, 1978.

[19] S. Dubuc. Interpolation through an iterative scheme. J. Math. Anal. Appl., 114:185–204, 1986.

[20] N. Dyn, J. Gregory, and D. Levin. Analysis of uniform binary subdivision schemes for curve design. Constr. Approx., 7:127–147, 1991.

[21] N. Dyn, D. Levin, and J. Gregory. A 4-point interpolatory subdivision scheme for curve design. Comput. Aided Geom. Des., 4:257–268, 1987.

[22] N. Dyn, D. Levin, and J. A. Gregory. A butterfly subdivision scheme for surface interpolation with tension control. ACM Trans. on Graphics, 9(2):160–169, April 1990.

[23] N. Dyn, D. Levin, and C. A. Micchelli. Using parameters to increase smoothness of curves and surfaces generated by subdivision. Comput. Aided Geom. Des., 7:129–140, 1990.

[24] T. Eirola. Sobolev characterization of solutions of dilation equations. SIAM J. Math. Anal., 23(4):1015–1030, 1992.

[25] I. Guskov. Multivariate subdivision schemes and divided differences. Technical report, Department of Mathematics, Princeton University, 1998.

[26] I. Guskov, W. Sweldens, and P. Schröder. Multiresolution signal processing for meshes. In Computer Graphics Proceedings (SIGGRAPH '99), pages 325-334, 1999.

[27] P. S. Heckbert and M. Garland. Survey of polygonal surface simplification algorithms. Technical report, Carnegie Mellon University, 1997.

[28] L. Kobbelt. Interpolatory subdivision on open quadrilateral nets with arbitrary topology. In Proceedings of Eurographics 96, Computer Graphics Forum, pages 409–420, 1996.

[29] L. Kobbelt. Discrete fairing. In Proceedings of the Seventh IMA Conference on the Mathematics of Surfaces, pages 101–131, 1997.

[30] L. Kobbelt and P. Schröder. Constructing variationally optimal curves through subdivion. Technical Report CS-TR-97-05, Department of Computer Science, California Institute of Technology, 1997.

[31] J. M. Lane and R. F. Riesenfeld. A theoretical development for the computer generation of piecewise polynomial surfaces. IEEE Trans. Patt. Anal. Mach. Intell., 3(1):35–46, 1980.

[32] C. Loop. Smooth subdivision surfaces based on triangles. Master's thesis, University of Utah, Department of Mathematics, 1987.

[33] S. G. Mallat. Multiresolution approximations and wavelet orthonormal bases of $L^2(\mathbf{R})$. Trans. Amer. Math. Soc., 315(1):69–87, 1989.

[34] C. A. Micchelli and H. Prautzsch. Computing surfaces invariant under subdivision. Computer Aided Geometric Design, 4(4):321–328, 1987.

[35] U. Reif. A unified approach to subdivision algorithms near extraordinary vertices. Computer Aided Geometric Design, 12:153–174, 1995.

[36] O. Rioul. Simple regularity criteria for subdivision schemes. SIAM J. Math. Anal., 23(6):1544–1576, 1992.

[37] P. Schröder and D. Zorin, editors. Course Notes: Subdivision for Modeling and Animation. ACM SIGGRAPH, 1998.

[38] J. E. Schweitzer. Analysis and Application of Subdivision Surfaces. PhD thesis, University of Washington, 1996.

[39] W. Sweldens. The lifting scheme: A construction of second generation wavelets. SIAM J. Math. Anal., 29(2):511–546, 1997.

[40] W. Sweldens and P. Schröder. Building your own wavelets at home. In Wavelets in Computer Graphics, pages 15–87. ACM SIGGRAPH Course notes, 1996.

[41] L. F. Villemoes. Wavelet analysis of refinement equations. SIAM J. Math. Anal., 25(5):1433–1460, 1994.

[42] J. Warren. Subdivision methods for geometric design. Unpublished manuscript, Rice University `http://www.cs.rice.edu/~jwarren`.

[43] J. Warren. Binary subdivision schemes for functions over irregular knot sequences. In M. Daehlen, T. Lyche, and L. Schumaker, editors, Mathematical Methods in CAGD III. Academic Press, 1995.

[44] D. Zorin. C^k continuity of subdivision surfaces. Technical report, California Institute of Technology, 1996.

[45] D. Zorin, P.Schröder, and W. Sweldens. Interpolating subdivision for meshes with arbitrary topology. In Computer Graphics Proceedings (SIGGRAPH '96), pages 189–192, 1996.