Hierarchical and Variational Geometric Modeling with Wavelets

Citation

Gortler, Steven J. and Michael F. Cohen. 1995. Hierarchical and variational geometric modeling with wavelets. In Proceedings of the 1995 symposium on interactive 3D graphics (SIGGRAPH 1995), April 9-12, 1995, Monterey, California, ed. SIGGRAPH and Michael J. Zyda, 35-42, 205. New York: ACM Press.

Published Version

<http://dx.doi.org/10.1145/199404.199410>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:2634298>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Hierarchical and Variational
Geometric Modeling with Wavelets

Steven J. Gortler and Michael F. Cohen

Work completed in part at
Department of Computer Science
Princeton University

July 20, 1995

Technical Report
MSR-TR-95-25

Microsoft Research
Advanced Technology Division
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

© ACM, 1995. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in the Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics, 35–42, 205. <http://doi.acm.org/10.1145/199404.199410>

Hierarchical and Variational Geometric Modeling with Wavelets

Steven J. Gortler and Michael F. Cohen

This paper also appears in the proceedings of the 1995 Symposium on Interactive 3D Graphics, April 1995, Monterey, CA.

Abstract

This paper discusses how wavelet techniques may be applied to a variety of geometric modeling tools. In particular, wavelet decompositions are shown to be useful for hierarchical control point or least squares editing. In addition, direct curve and surface manipulation methods using an underlying geometric variational principle can be solved more efficiently by using a wavelet basis. Because the wavelet basis is hierarchical, iterative solution methods converge rapidly. Also, since the wavelet coefficients indicate the degree of detail in the solution, the number of basis functions needed to express the variational minimum can be reduced, avoiding unnecessary computation. An implementation of a curve and surface modeler based on these ideas is discussed and experimental results are reported.

1 Introduction

Wavelet analysis provides a set of tools for representing functions hierarchically. These tools can be used to facilitate a number of geometric modeling operations easily and efficiently. In particular, this paper explores three paradigms for free-form curve and surface construction: control point editing, direct manipulation using least squares, and direct manipulation using variational minimization techniques. For each of these paradigms, the hierarchical nature of wavelet analysis can be used to either provide a more intuitive modeling interface or to provide more efficient numerical solutions.

In control point editing, the user sculpts a free-form curve or surface by dragging a set of control points. A better interface allows the user to directly manipulate the curve or surface itself, which defines a set of constraints. In a *least squares* paradigm, given a current curve or surface, the modeling tool returns the curve or surface that meets the constraints by changing the current control points by the least squares amount [1, 11].

The behavior of the modeling tool is determined by the type of control points and *basis functions* used to describe the curve or surface. With the uniform cubic B-spline basis, for example, the user's actions result in local changes at a predetermined scale. This is not fully desirable; at times the user may want to make fine changes of detail, while at other times he may want to easily make broad changes. Hierarchical B-splines offer a representation that allows both control point and least squares editing to be done at multiple resolutions [9]. Hierarchical B-splines, though, form an over-representation for curves and surface (i.e., any curve has multiple representations using hierarchical B-splines). As a result, the same curve may behave differently to a user depending on the particular underlying representation. In contrast, B-spline wavelets form a hierarchical basis for the space of B-spline curves and surfaces in which every object has a unique representation. Wavelet methods in conjunction with hierarchical B-splines provide a method for constructing a useful geometric modeling interface. This approach is similar to the one described by Finkelstein and Salesin [8]. In this paper we will discuss some of the various issues that are relevant to building such a modeling tool.

Variational modeling is a third general paradigm for geometric modeling [2, 28, 21]. In this setting, a user alters a curve or surface by directly manipulation, as above, defining a set of constraints. The variational modeling paradigm seeks the "best" solution amongst all answers that meet the constraints. The notion of best, which is formally defined as the solution that *minimizes some energy function*, is often taken to mean the *smoothest* solution.

In theory, the desired solution is the curve or surface that has the minimum energy of *all* possible curves or surfaces that meet the constraints. Unfortunately there is little hope to find a closed form solution¹. Therefore, in practice, the "space" of parametric curves or surfaces is restricted to those represented by a linear combination of a fixed set of basis functions such as cubic B-splines. Given a set of n basis functions, the goal of finding the best curve or surface is then reduced to that of finding the best set of n coefficients. This reduction is referred to as the *finite element method* [27].

The general case requires solving a non-linear optimization problem. In the best case, the energy function is quadratic and the constraints are linear leading to a single linear system to solve. But even this can be costly when n is large since direct methods for matrix inversion require $O(n^3)$ time. To accelerate this process it is tempting to use gradient-type iterative methods to solve the linear system; these methods only take $O(n)$ time per iteration, due to the $O(n)$ matrix sparsity created by the finite element formulation. Unfortunately, the linear systems arising from a finite element formulation are often expensive to solve using iterative methods. This is because the systems are ill-conditioned, and thus require many iterations to converge to a minimum [26, 25]. Intuitively speaking this occurs because each basis function represents a very narrow region of the answer; there is no basis function which can be moved to change the answer in some broad manner. For example, changing one coefficient in a cubic B-spline curve during an iteration alters the curvature in a local region only. In order to produce a broad smooth curve, the coefficients of the neighboring B-splines will move in next few iterations. Over the next many iterations, the solution process will affect wider and wider regions, and the effect will spread out slowly like a wave moving along a string. The result is very slow convergence (see Figure (1)). One method used to combat this problem is multigriding [26, 10], where a sequence of problems at different resolution levels are posed and solved.

An alternative approach, is to use a *wavelet* basis instead of a standard finite element basis [25, 23, 15, 22]. In a wavelet basis, the answer is represented hierarchically. This allows the solution method to alter the answer at any desired resolution by altering the proper basis function, and thus the ill-conditioning is avoided. In this paper we show how to use a wavelet construction, which is based on cubic B-splines, to quickly solve variational modeling problems in an elegant fashion.

Another problem with the finite element approach is choosing the density of the basis functions. If too few basis functions (too few B-spline segments or tensor product B-spline patches) are used then the solution obtained will be far from the actual minimum. If too many basis functions are used then unnecessary computation will be performed during each iteration (n is too big). In order to successfully choose a proper density, one must know how much detail exists in the variational minimum answer. Since, a priori, this is unknown, an efficient solver must be able to adaptively change the basis during the solution process [28], one needs an easy way to detect that too many or too few basis functions are being used. In addition, one needs a basis for which adding more detail, (i.e., refinement), is easy. Wavelets offer a basis where this task can be accomplished quickly and elegantly.

The work presented in this paper combines the wavelet approaches of [25], [12], and [16]. Like [25], this paper uses hierarchical basis functions as a pre-conditioner, so that fewer iterations

¹But see [20].

Figure 1: Minimum energy solutions subject to three constraints, found by the B-spline and wavelet methods after various numbers (0-1024) of iterations. (65 variables, 3 constraints). This illustrates the ill conditioning of the B-spline optimization problem.

are needed for convergence. Similar to [12] and [16], wavelets are also used as a method for limiting the solution method to the proper level of detail.

2 Geometric Representation

This paper will restrict itself to parametric representations of curves and surfaces. In this representation, a curve is defined as a 3 dimensional trajectory parameterized by t ,

$$\gamma(t) = (X(t), Y(t), Z(t)) \quad (1)$$

and a surface is defined as

$$\gamma(s, t) = (X(s, t), Y(s, t), Z(s, t)) \quad (2)$$

which defines a three dimensional location for every parameter pair (s, t) .

The parametric representation of a curve or surface is made up of three functions X, Y, Z , which are represented as a linear combination of basis functions. Just focusing on the X function, for curves this becomes

$$X(t) = \sum_j x_j \phi_{L,j}(t) \quad (3)$$

and for surfaces

$$X(s, t) = \sum_{j,k} x_{j,k} \phi_{L,j,k}(s, t) \quad (4)$$

where the x are scalar coefficients. In geometric modeling the univariate basis $\phi_{L,j}(t)$ is typically some “piecewise” basis, such as a cubic B-spline or the Bernstein (Bézier) basis, and the bivariate basis used for surfaces is the associated tensor product basis $\phi_{L,j,k}(s, t) \equiv \phi_{L,j}(s)\phi_{L,k}(t)$.

3 Hierarchical Geometric Descriptions

In this section we will briefly review some ways that curves and surfaces may be represented hierarchically.

Let us begin by discussing curves. For simplicity we will deal with the uniform cubic B-spline basis over the interval $[0 \dots 2^L]$ made up of translations of a single basis shape denoted $\phi(t)$. The cubic B-spline function $\phi(t)$ is supported over the interval $[0 \dots 4]$ and is made up of 4 cubic polynomial pieces joined with C^2 continuity. The complete uniform cubic B-spline basis is made up of translated copies $\phi_{L,j}(t)$ of the basis shape $\phi(t)$ (see Figure 2).

$$\phi_{L,j}(t) = \phi(t - j) \quad (5)$$

The index j represents the translation of a specific basis from the canonical B-spline left justified at zero, and L is the *level* or resolution of the basis. There are roughly 2^L functions in this basis². In wavelet terminology, the space (or family) of curves spanned by all linear combinations of these basis functions is denoted V_L (e.g., V_L contains all functions that are piecewise cubic, with simple knots at the integers).

3.1 Hierarchical B-splines

Forsey and Bartels [9] introduced hierarchical B-splines as a way of representing and modeling geometric objects hierarchically. Instead of using only B-spline basis functions at a single resolution L , they use a hierarchy of wider and wider B-spline functions

$$\phi_{i,j}(t) = \phi(2^{L-i}t - j) \quad (6)$$

for $0 \leq i \leq L$. For example, the basis functions $\phi_{L-1,j}$ at resolution level $L - 1$ (with a support size of 8), are twice as wide as the basis functions $\phi_{L,j}$ at level L (with a support size of 4). These basis functions, $\phi_{L-1,j}$, span the space of piecewise cubic functions with knots at all *even* integers; in wavelet terminology, this space is called V_{L-1} . On each coarser level, the space V_i has half as many basis functions, and they are all twice as wide.

According to the well known B-spline knot insertion algorithm [6, 9, 3] one can define double width B-spline basis functions as linear combinations of single width B-spline basis functions.

$$\phi_{i-1,j} = \sum_k h_{k-2j} \phi_{i,k} \quad (7)$$

where the sequence h is

$$h[0..4] = \left\{ \frac{1}{8}, \frac{4}{8}, \frac{6}{8}, \frac{4}{8}, \frac{1}{8} \right\} \quad (8)$$

(see Figure (2)). As a result of Equation (7) the set of functions in V_{i-1} is a subset of the functions in V_i .

$$V_{i-1} \subset V_i \quad (9)$$

The basic idea of Forsey and Bartels is to allow the user to control the coefficient of each of these basis functions $\phi_{i,j}$ by exposing a control mesh at each level i .

3.2 Wavelets

Hierarchical B-splines $\{\phi_{i,j}\}$ do not form a *basis* for the function space V_L ; they form an *overrepresentation* for all the curves in

² A few extra basis functions are needed at the boundary. This paper will not discuss the technical details needed to handle all of the boundary constraints. This is discussed in many places including [4, 16, 8, 13].

V_L . In other words, there are many linear combinations of the basis functions defining the same curve or surface. Wavelets are a representation related to hierarchical B-splines, that form a basis; in a wavelet basis, all curves in V_L have a unique representation.

Rather than add a new finer set of B-splines at each level of the hierarchy, the idea is to look for a set of functions $\psi_{i,j}$ that “fills in” the space between the adjacent B-spline spaces, V_i and V_{i+1} . These wavelet functions $\psi_{i,j}$ represent the *detail* of the curve that cannot be represented by the double width B-splines, $\phi_{i,j}$. For each i , the space of functions spanned by the $\psi_{i,j}$ is called W_i .

There is actually quite a bit of freedom in choosing these $\psi_{i,j}$ functions, and hence the space W_i , as long as every function in V_{i+1} can be written as a combination of some function in V_i and some function in W_i . This is notated as

$$V_{i+1} = V_i + W_i \quad (10)$$

Just like the Hierarchical B-splines are all scales and translates of a single shape $\phi(t)$, (see Equation (5)) in a wavelet basis, the basis functions $\psi_{i,j}$ are all translates and scales of a single function $\psi(t)$.

$$\psi_{i,j}(t) = \psi(2^{L-i}t - j) \quad (11)$$

Also similar to hierarchical B-splines, in a wavelet basis, the basis functions on one level can be defined by linearly combining B-spline functions on the next finer resolution,

$$\psi_{i-1,j} = \sum_k g_{k-2j} \phi_{i,k} \quad (12)$$

And as a result $W_{i-1} \subset V_i$. There is some degree of freedom in choosing the sequence g , as long as the property expressed by Equation (10) holds. One such sequence given by Cohen et al. [5] is ³ (see Figure (3)).

$$g[0..10] = \left\{ \frac{5}{256}, \frac{20}{256}, \frac{1}{256}, \frac{-96}{256}, \frac{-70}{256}, \frac{280}{256}, \frac{-70}{256}, \frac{-96}{256}, \frac{1}{256}, \frac{20}{256}, \frac{5}{256} \right\}$$

Due to the relationships of Equations (7) and (12), if some function $X(t)$ in V_i has been expressed as a linear combination of the B-spline basis function at level $i-1$ and wavelet basis functions at level $i-1$, using coefficients notated by $x_{\phi_{i-1,j}}$ and $x_{\psi_{i-1,j}}$,

$$X(t) = \sum_j x_{\phi_{i-1,j}} \phi_{i-1,j}(t) + x_{\psi_{i-1,j}} \psi_{i-1,j}(t) \quad (13)$$

then, $x_{\phi_{i,j}}$, the coefficients of the same function, with respect to the B-spline basis at level i may be found with

$$x_{\phi_{i,j}} = \sum_k h_{j-2k} x_{\phi_{i-1,k}} + \sum_k g_{j-2k} x_{\psi_{i-1,k}} \quad (14)$$

and now $X(t) = \sum_j x_{\phi_{i,j}} \phi_{i,j}(t)$

Inversely, if some function has been expressed with respect to B-spline functions at level i , then the representation of Equation (13) may be found using the formula

$$x_{\phi_{i-1,j}} = \sum_k \tilde{h}_{k-2j} x_{\phi_{i,k}} \quad (15)$$

$$x_{\psi_{i-1,j}} = \sum_k \tilde{g}_{k-2j} x_{\phi_{i,k}} \quad (16)$$

using the proper inverse sequences $\tilde{g}$ and $\tilde{h}$. Equation (15) *projects* the high resolution curve from V_i into the lower resolution space V_{i-1} ; this is, in some sense, a smoother approximation of the object

³A different sequence is given by Chui [3] and generates a semi-orthogonal wavelet.

Figure 2: Five B-splines $\phi_{L,j}$ may be combined using the weights h to construct the double width B-spline $\phi_{L-1,0}$

Figure 3: Eleven B-splines $\phi_{L,j}$ may be combined using the weights g to construct the wavelet function $\psi_{L-1,0}$

in V_i . Equation (16) captures the detail that is lost in this projection, and represents it using a basis for the space W_{i-1} .

When using the h and g sequences given by Cohen et al [5], the proper inverse sequences $\tilde{h}$ and $\tilde{g}$ are

$$\begin{aligned} \tilde{h}[-3..7] &= \left\{ \frac{-5}{256}, \frac{20}{256}, \frac{-1}{256}, \frac{-96}{256}, \frac{70}{256}, \frac{280}{256}, \frac{70}{256}, \frac{-96}{256}, \frac{-1}{256}, \frac{20}{256}, \frac{-5}{256} \right\} \\ \tilde{g}[3..7] &= \left\{ \frac{1}{8}, \frac{-4}{8}, \frac{6}{8}, \frac{-4}{8}, \frac{1}{8} \right\} \end{aligned} \quad (17)$$

3.3 The Basis

Every function in V_L , expressed as a combination of the B-spline basis functions $\{\phi_{L,j}\}$, can be expressed uniquely in the *wavelet basis* is made up by the functions

$$\{\phi_{0,j}, \psi_{i,j}\} \quad 0 \leq i \leq L-1 \quad (18)$$

In the wavelet representation, the function is expressed hierarchically.

Transforming a function's representation from B-spline to wavelet coefficients may be done with the pyramid procedure **coef_pyrm_up**. This procedure may be performed in linear time by successively applying the transformation of Equations (15) and (16). This linear transformation may be denoted by the matrix $\mathbf{W}$. The inverse transformation (denoted by the matrix $\mathbf{W}^{-1}$), may be implemented with the procedure **coef_pyrm_down**, which successively applies the transformation of Equation (14).

If **coef_pyrm_up** is implemented using the h and g sequences instead of the $\tilde{h}$ and $\tilde{g}$ sequences, then the resulting procedure may be called **basis_pyrm_up**, and it is represented by the matrix $\mathbf{W}^{-\mathbf{T}}$. If **coef_pyrm_down** is implemented using the $\tilde{h}$ and $\tilde{g}$ sequences instead of the h and g sequences, then the resulting procedure may be called **basis_pyrm_down**, and it is represented by the matrix $\mathbf{W}^{\mathbf{T}}$.

3.4 Surfaces

The ideas outlined above are easily extended to tensor product surfaces [3]. The uniform tensor product cubic B-spline basis is made

up of the functions $\phi_{L,j}(s)\phi_{L,k}(t)$. The hierarchical uniform tensor product cubic B-spline representation is made up of the functions $\phi_{i,j}(s)\phi_{i,k}(t)$ for $0 \leq i \leq L$. On each coarser resolution of the hierarchy, there are $1/4$ the amount of ϕ basis functions.

The tensor product B-spline wavelet basis is made up of the functions ⁴

$$\begin{matrix} \phi_{0,j}(s)\phi_{0,k}(t) & \phi_{i,j}(s)\psi_{i,k}(t) \\ \psi_{i,j}(s)\phi_{i,k}(t) & \psi_{i,j}(s)\psi_{i,k}(t) \end{matrix} \quad (19)$$

with i in $\{0 \dots L-1\}$.

Just like for curves, there are four pyramid procedures and associated **W** matrices.

4 Geometric Modeling with Wavelets

The styles of interactive control discussed in the introduction will be revisited in the context of hierarchical representations. *Multiresolution modeling* allows the user to interactively modify the curve or surface at different resolution levels. This allows the user to make broad changes while maintaining the details, and conversely detailed changes while maintaining the overall shape. Two types of hierarchical manipulation are considered, control point dragging and a direct manipulation involving solving a least squares problem.

In contrast, *variational modeling* allows the user to directly manipulate the curve or surface with the curve or surface maintaining some notion of overall smoothness subject to user imposed constraints. This physically based paradigm provides an intuitive means for shape control. Each of these paradigms will be explored in the context of wavelet bases which will be shown to provide the required hooks for such interaction and/or significant computational savings.

4.1 Multiresolution Modeling

A multiresolution representation such as a hierarchical B-spline or wavelet representation may be used to implement a multiresolution modeling system. This section explores the choices that must be made when designing a multiresolution tool. Two related methods are described; direct control point manipulation and a least squares solver.

In control point modeling, the user is allowed to directly alter the coefficient values, by clicking and dragging on control points. In the least squares scheme [1, 11], the user can click and drag directly on the curve or surface, defining interpolation and tangent constraints. The system returns the curve or surface that satisfies these linear constraints ($\mathbf{Ax} = \mathbf{b}$), by changing the coefficients by the least squares amount. Least square solutions can be found very inexpensively using the pseudoinverse [11]. The least squared problem can also be posed as a minimization problem [28], whose solution can be found by solving a sparse, well conditioned, linear system.

In multiresolution versions of these two schemes, the user chooses the resolution level i , and then only the quantities of basis functions on level i are altered. The locality of the effect on the curve or surface is directly tied to the chosen level i . In control point modeling, the control polygon at level i is manipulated by the user. In a least squares scheme, the user is provided a direct handle on the curve or surface itself, and the least squares solution is found only using the basis functions on level i . The least-squares approach offers a much more intuitive interface, and (for curves) works at interactive speeds.

One decision to be made is whether to expose the user to hierarchical B-splines or to wavelets. It is easy to see that manipulating wavelet basis functions does not produce an intuitive interface. Moving such a control point, and thus changing the amount of some

Figure 4: When B-spline coefficients are manipulated, the curve responds in a “hump” like fashion. When wavelet coefficients are manipulated, the curve responds in a “wave” like fashion.

Figure 5: When the (X,Y,Z) frame is used for wavelet multiresolution editing, detail maintains its orientation as the sweep is changed. When the normal, tangent, bi-normal, (N, T, B) frame is used with a wavelet representation, the detail does not maintain its structure as the sweep is changed. When the (N, T, B) frame is used with a B-spline representation, the detail follows the orientation of the curve.

wavelet basis function used, changes the solution in a “wave” like fashion. In contrast, it is more intuitive to move a B-spline control point which changes the solution in a “hump” like fashion (see Figure 4). Thus the user in this case should manipulate the hierarchical B-spline functions.

4.2 Orientation

In the parametric representation, the curve or surface is represented by three functions X, Y, Z . In the the multi-resolution paradigm, when a user adds fine directional detail, say a fine hump in the X direction, this detail will become locked in the originally chosen direction. If the user later manipulates the broad sweep of the curve, the detail will maintain its original direction (see Figure 5). This is not always desirable, since the user may want the detail’s orientation to follow the changing direction of broader curve or surface.

An “orientation” approach first proposed by Forsey and Bartels [9] may be applied to the multiresolution editing scheme. In a multiresolution modeling system all of the information describing the curve or surface lives at some resolution. In an orientation approach, the information at each resolution i is not expressed as three independent functions of (X, Y, Z) . Instead the detail at each resolution i is represented with respect to the geometric shape of the lower resolution version of the curve or surface. This lower resolution version is defined by summing all of the information from all the lower resolution levels.

Tangent and normal directions of the lower resolution curve or surface are then computed at a series of sample points. The detail coefficients at level i are then expressed with respect to these tangent and normal directions instead of the (X, Y, Z) directions. If any

⁴This basis is known as the non-standard basis [3].

lower resolution component of the curve is later explicitly altered, then the detail's orientation will change appropriately.

4.2.1 Defining Detail

In order to apply an orientation approach, one must have some method for decomposing the object into components at different resolutions. When one is using hierarchical B-splines, which over-represent objects in V_L , then there is some freedom in defining what information resides at which level of detail.

If the geometric object is being designed with a multiresolution editor, then the user is explicitly manipulating the object at resolutions that he chooses. Therefore, one simple method is to maintain all information at the resolution entered by the user [9]. Using this method, the same geometric object may behave differently depending on the way the object was generated.

An alternative is to use wavelet analysis: begin with the complete resolution object (in V_L), and then successively *project* it to each lower resolution level using Equation (15). This generates a unique smoothed version of the object at each resolution V_i . The object can now be represented as a combination of components from the difference spaces W_i .

In typical wavelet analysis, the components in W_i are represented using some special basis functions $\psi_{i,j}$ that span the difference space W_i . Alternatively, instead of using wavelet functions $\psi_{i,j}$ to represent the difference, one may instead use the B-spline functions on the next finer level $\phi_{i+1,j}$. This can be done because of Equation (12). The choice of whether to use B-spline or wavelets to represent the functions in W_i is an important question that we shall deal with soon.

4.2.2 Projections between Levels

There are many ways to obtain a lower resolution version of some object from V_L . For example, given an object in V_L , one could obtain a lower resolution version in V_{L-1} by throwing away every other control point. Subsampling is not a true projection; starting with a smooth curve in V_{L-1} , and then expressing that smooth curve in the higher resolution B-spline basis V_L , and finally subsampling the control points will not return the original smooth curve we began with.

Another way of obtaining a smoothed version of the object is by *orthogonally* projecting the object from V_L into V_{L-1} . The orthogonal projection is the object in V_{L-1} that is closest to object in V_L using the L^2 measure. One may obtain the orthogonal projection by using Equation (15), with the $\bar{h}$ sequence given for the semi-orthogonal wavelet construction by Chui [3]. This is the approach used in [8]. Although this is a very elegant way of obtaining a lower resolution version of an object, it has a few drawbacks. This particular $\bar{h}$ sequence is infinite in length (although it does decay rapidly from its centers) and so performing this task efficiently can be troublesome. Also, because these sequences are not local, then a single change to one B-spline coefficient at level L will alter all of the coefficients of the projection at level $L - 1$.

One good compromise between these two extremes (subsampling, and orthogonal projection), is to use Equation (15) but to use the $\bar{h}$ filter given for the non-orthogonal wavelet construction by Cohen et al. [5]. This projection is non-orthogonal, but it is entirely local. This is the choice we have used in our multiresolution modeling tool.

4.2.3 Representing Detail

What set of basis functions should be used to represent the detail. If a wavelet projection Equation (15) is used to define the lower resolution versions of the object, then the detail can be represented

by using the corresponding wavelet functions. The other option is to represent the detail using hierarchical B-spline functions. The disadvantage of using hierarchical B-splines is that there are roughly $2n$ B-splines in the hierarchy, and only n wavelets.

The advantage of using hierarchical B-splines however is that they maintain the orientation better. When the user changes the broad sweep of the curve, changing the tangent, normal, and bi-normal frame at t_j , the detail functions are remixed. If the detail functions are wavelet functions, then changing the normal and tangent frame remixes "wave" shaped functions introducing non-intuitive wiggles. If the detail functions are B-spline basis functions, then "hump" shaped functions get remixed, yielding more intuitive changes. Also if the detail functions are B-splines, then because there are twice as many B-splines than wavelets, the tangent and normal directions are computed at twice as many sample points allowing the detail to follow the orientation with more fidelity (see Figure 5).

5 Variational Modeling

The variational modeling paradigm generalizes the least squares notion to any *objective* function minimization, typically one representing minimizing curvature. The variational problem leads to a non-linear optimization problem over a finite set of variables when cast into a given basis.

There are a variety of objective functions used in geometric modeling [21, 24] In our implementation we have used the *thin-plate* measure which is based on parametric second derivatives [27, 2, 28]. The thin plate minimum may be found by solving the following linear system [28].

$$\begin{bmatrix} \mathbf{H} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{b} \end{bmatrix} \quad (20)$$

Where $\mathbf{A}$ is the constraint matrix, $\mathbf{H}$ is the Hessian matrix, and λ are Lagrange variables.

5.1 Hierarchical Conditioning

Wavelets can be used in the context of variational modeling so that the solution may be obtained more efficiently.

In the B-spline basis, the optimization procedure resulted in the linear system given by Equation (20). In the wavelet basis, a different linear system results which is given by

$$\begin{bmatrix} \bar{\mathbf{H}} & \bar{\mathbf{A}}^T \\ \bar{\mathbf{A}} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{b} \end{bmatrix} \quad (21)$$

where the bars signify that the variables are wavelet coefficients, $\bar{\mathbf{x}} = \mathbf{W}\mathbf{x}$, and the Hessian and constraint matrix are expressed with respect to the wavelet basis. To see the relationship with the B-spline system, the new system can also be written down as

$$\begin{bmatrix} \mathbf{W}^{-T}\mathbf{H}\mathbf{W}^{-1} & \mathbf{W}^{-T}\mathbf{A}^T \\ \mathbf{A}\mathbf{W}^{-1} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \bar{\mathbf{x}} \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{b} \end{bmatrix} \quad (22)$$

Although Equation (20) and Equation (21/22) imply each other, they are two distinct linear systems of equations. Because the wavelet system (21/22) is hierarchical it will not suffer from the poor conditioning of the B-spline system of Equation (20). For a rigorous discussion of the relevant theory see [7].

The scaling of the basis functions is very significant for the behavior of the optimizing procedures. Traditionally the wavelet functions are defined with the following scaling [19, 22]:

$$\begin{aligned}\phi_{i,j}(t) &= 2^{(i-L)/2} \phi(2^{(i-L)}t - j) \\ \psi_{i,j}(t) &= 2^{(i-L)/2} \psi(2^{(i-L)}t - j)\end{aligned}\quad (23)$$

This means that at each level moving up, the basis functions become twice as wide, and are scaled $\frac{1}{\sqrt{2}}$ times as tall. While in many contexts this normalizing may be desirable, for optimization purposes it is counter productive. For the optimization procedure to be well conditioned [15, 7] it is essential to emphasize the coarser levels. The correct theoretical scaling depends on both the energy function used, and the dimension of problem. For a fuller discussion, see the Appendix in [13]. In the experiments described in this paper the following scaling was used

$$\begin{aligned}\phi_{i,j}(t) &= 2^{-(i-L)} \phi(2^{(i-L)}t - j) \\ \psi_{i,j}(t) &= 2^{-(i-L)} \psi(2^{(i-L)}t - j)\end{aligned}\quad (24)$$

This means that as one goes from level i to level $i - 1$ the basis functions become twice as wide, and $1/2$ as tall. In the pyramid code, this is achieved by multiplying all of the h and g entries by 2, and all of the $\tilde{h}$ and $\tilde{g}$ by $1/2$ ⁵.

5.1.1 Explicit vs. Implicit

There is now a choice to make. In an iterative conjugate gradient solver, the common operation is multiplication of a vector times the wavelet matrix given in Equations (21/22). There are two ways to implement this.

One approach, the *explicit* approach, is to compute and store the wavelet Hessian matrix $\tilde{\mathbf{H}}$ and the wavelet constraint matrix $\tilde{\mathbf{A}}$ (Equation (21)). These can be computed directly from a closed form (piecewise polynomial) representation of the wavelet functions $\psi_{i,j}$. Unfortunately, these matrices are not as sparse as the B-spline Hessian and constraint matrices.

Alternatively, there is the *implicit* approach [29, 25] which only computes and stores the entries of the B-spline matrices $\mathbf{H}$ and $\mathbf{A}$ (Equation (22)). Multiplication by the $\mathbf{W}$ matrices is accomplished using the **pyrm** procedures. The advantage of this approach is that the whole multiply remains $O(n)$ in both time and space, since the **pyrm** procedures run in linear time, and the matrices $\mathbf{H}$ and $\mathbf{A}$ are $O(n)$ sparse. Even though one of the methods explicitly uses wavelet terms while the other uses B-spline terms, these two methods are mathematically equivalent, and so both will have the same condition properties.

5.2 Adaptive Oracle

By limiting the possible surfaces to only those that can be expressed as a linear combination of a fixed set of basis functions, one obtains an approximation of the true optimal surface. As more basis functions are added, the space of possible solutions becomes richer and a closer approximation to the true optimal surface can be made. Unfortunately, as the space becomes richer, the number of unknown coefficients increases, and thus the amount of computation required per iteration grows. A priori, it is unknown how many basis functions are needed. Thus, it is desirable to have a solution method that adaptively chooses the appropriate basis functions. This approach was applied using hierarchical B-splines in [28]. When refinement was necessary, “thinner” B-splines basis functions were added, and the redundant original “wider” B-splines

were removed. With wavelets, all that must be done is to add in new “thinner” wavelets wherever refinement is deemed necessary. Since the wavelets coefficients correspond directly to local detail, all previously computed coefficients are still valid.

The decision process of what particular wavelets to add and remove is governed by an **oracle** procedure which is called after every fixed number of iterations. The oracle must decide what level of detail is required in each region of the curve or surface.

When some region of the solution does not need fine detail, the corresponding wavelet coefficients are near zero, and so the first thing the **oracle** does is to deactivate the wavelet basis functions whose corresponding coefficients are below some small threshold. The **oracle** then activates new wavelet basis functions where it feels more detail may be needed. There are two criteria used. If a constraint is not being met, then the oracle adds in finer wavelet functions in the region that is closest in parameter space to the unmet constraint. Even if all the constraints are being met, it is possible that more basis functions would allow the freedom to find a solution with lower energy. This is accomplished by activating finer basis functions near those with coefficients above some maximum threshold.

To avoid cycles, a basis function is marked as being **dormant** when it is removed from consideration. Of course, it is possible that later on the solution may really need this basis function, and so periodically there is a **revival** phase, where the **dormant** marks are removed.

5.3 User Interface

A user of the system is first presented with a default curve or surface. Constraints can then be introduced by clicking on the curve or surface with the mouse. The location of the mouse click defines a parametric position t (and s) on the curve (or surface). The user can then drag this point to a new location to define an interpolation constraint. Tangent constraints at a point can also be defined by orienting “arrow” icons at the point. Once the constraint is set, the solver is called to compute the minimum energy solution that satisfies the constraints placed so far. Resulting curves and surfaces are displayed using SGI GL **nurbscurve** and **nurbssurface** calls⁶.

When the solution is completed, the result provides information for not only the curve or surface satisfying the specific value of the new constraint, but for all curves or surfaces with respect to any value of this constraint. Once the linear system (Equation (21/22)) with the newest constraint has been solved, the solver stores the delta vector

$$\frac{\Delta \bar{x}}{\Delta b_m} \quad (25)$$

where m is the index of the newest constraint, and b_m is the constraint value (i.e., the position or tangent specified by the user). This vector stores the change of the coefficient vector due to a unit change in the new constraint Δb_m , essentially a column of the inverse matrix. The user is now free to interactively move the target location of the constraint without having to resolve the system since, as long as the parameters s , and t of the constraints do not change, the matrix of the system, and thus its inverse, do not change. However, as soon as a new constraint is added (or a change to the parameters s and t is made) there is fresh linear system that must be solved, and all of the delta vectors are invalidated. The ability to interactively change the value of a constraint is indicated to the user by coloring the constraint icon. See Color Plate.

⁵The proper scaling is essential to obtain the quick convergence of the wavelet method when steepest descent or conjugate gradient iteration is used. Scaling is not important with Gauss-Seidel iteration, which will perform the same sequence of iterations regardless of scale.

⁶One GL call to **nurbssurface** can be more expensive than a complete iteration.

Figure 6: Error per time. Curve with 65 control points, 3, 7, and 13 constraints.

5.4 Variational Modeling Results

A series of experiments were conducted to examine the performance of the wavelet based system compared to a B-spline basis. In the curve experiments, the number of levels of the hierarchy, L , was fixed to 6, and in the surface experiments, L was fixed as 5. The optimization process was then run on problems with different numbers of constraints. The results of these tests are shown in Figures 6 and 7. These graphs show the convergence behavior of three different methods, solving with the complete B-spline basis, solving with the complete wavelet basis, and solving with an adaptive wavelet basis that uses an oracle. (The wavelet results shown here are using the *implicit* implementation). If $\mathbf{x}^{(m)}$ is the computed solution expressed as B-spline coefficients at time m , and $\mathbf{x}^*$ is the correct solution of the complete linear system⁷ (i.e., the complete system with $2^L + 1$ variables, and no adaptive oracle being used) then the error at time m is defined as

$$\frac{\sum_j |x_j^* - x_j^{(m)}|}{\sum_j |x_j^* - x_j^{(0)}|} \quad (26)$$

To obtain the starting condition $\mathbf{x}^{(0)}$, two constraints were initialized at the ends of the curve, and the minimal thin plate solution (which in this case is a straight line) was computed. (For surfaces, the four corners were constrained.) All times were taken from runs on an SGI R4000 reality engine.⁸

When there are large gaps between the constraints, the B-spline method is very poorly conditioned, and converges quite slowly while the wavelet method converges dramatically faster. In these problems, the oracle decides that it needs only a very small active set of wavelets and so the adaptive method converges even faster. As the number of constraints is increased, the solution becomes more tightly constrained, and the condition of the B-spline system

⁷computed numerically to high accuracy

⁸In the curve experiments, each B-spline iteration took 0.0035 seconds, while each iteration of the implicit wavelet method took 0.011 seconds. For the surface experiments, each B-spline iteration took 0.68 seconds while each iteration of the implicit wavelet method took 0.85 seconds. (The wavelet iterations using the explicit representation took about 10 times as long).

Figure 7: Error per time. Surface with 1089 control points, 11,23,64 evenly space constraints, and 62 constraints along the boundary.

improves. (Just by satisfying the constraints, the B-spline solution is very close to minimal energy). Meanwhile the oracle requires a larger active set of wavelets. Eventually, when enough constraints are present, the wavelet methods no longer offer an advantage over B-splines.

Experiments were also run where all the constraints were along the boundary of the surface. In these experiments there are many constraints, but the since the constraints are along the boundary, much of the surface is “distant” from any constraint. In these problems, the wavelets also performed much better than the B-spline method.

6 Conclusion

This paper has explored the use of wavelet analysis in a variety of modeling settings. It has shown how wavelets can be used to obtain multiresolution control point and least squares control. It has shown how wavelets can be used to solve variational problems more efficiently.

Future work will be required to explore the use of higher order functionals like those given in [21, 24]. Because the optimization problems resulting from those functionals are non-linear, they are much more computationally expensive, and it is even more important to find efficient methods. It is also important to study optimization modeling methods where constraint changes only have local effects.

Many of these concepts can be extended beyond the realm of tensor product uniform B-splines. Just as one can create a ladder of nested function spaces V_i satisfying the property of Equation (10) using uniform cubic B-splines of various resolutions, one can also create a nested ladder using non-uniform B-splines [18].

Subdivision surfaces are a powerful technique for describing surfaces with arbitrary topology [14]. A subdivision surface is defined by iteratively refining an input control mesh. As explained by Lounsbery et al. [17], one can develop a wavelet decomposition of such surfaces. Thus, many of the ideas developed in this paper may be applicable to that representation as well.

Acknowledgements

We are grateful to James H. Shaw for developing the graphical interface to the optimization program.

REFERENCES

- [1] BARTELS, R., AND BEATTY, J. A Technique for the Direct Manipulation of Spline Curves. In *Graphics Interface 1989* (1989), pp. 33–39.
- [2] CELNIKER, G., AND GOSSARD, D. Deformable Curve and Surface Finite-Elements for Free-From Shape Design. *Computer Graphics* 25, 4 (July 1991), 257–266.
- [3] CHUI, C. K. *An Introduction to Wavelets*, vol. 1 of *Wavelet Analysis and its Applications*. Academic Press Inc., 1992.
- [4] CHUI, C. K., AND QUAK, E. Wavelets on a Bounded Interval. *Numerical Methods of Approximation Theory* 9 (1992), 53–75.
- [5] COHEN, A., DAUBECHIES, I., AND FEAUVEAU, J. C. Biorthogonal Bases of Compactly Supported Wavelets. *Communication on Pure and Applied Mathematics* 45 (1992), 485–560.
- [6] COHEN, E., LYCHE, T., AND RIESENFELD, R. Discrete B-Splines and Subdivision Techniques in Computer-Aided Geometric Design and Computer Graphics. *Computer Graphics and Image Processing* 14, 2 (October 1980), 87–111.
- [7] DAHMEN, W., AND KUNOTH, A. Multilevel Preconditioning. *Numerische Mathematik* 63 (1992), 315–344.
- [8] FINKELSTEIN, A., AND SALESIN, D. Multiresolution Curves. In *Computer Graphics, Annual Conference Series, 1994* (1994), Siggraph, pp. 261–268.
- [9] FORSEY, D., AND BARTELS, R. Hierarchical B-Spline Refinement. *Computer Graphics* 22, 4 (August 1988), 205–212.
- [10] FORSEY, D., AND WENG, L. Multi-resolution Surface Approximation for Animation. In *Graphics Interface* (1993).
- [11] FOWLER, B. Geometric Manipulation of Tensor Product Surfaces. In *Proceedings, Symposium on Interactive 3D Graphics* (1992), pp. 101–108.
- [12] GORTLER, S., SCHRÖDER, P., COHEN, M., AND HANRAHAN, P. Wavelet Radiosity. In *Computer Graphics, Annual Conference Series, 1993* (1993), Siggraph, pp. 221–230.
- [13] GORTLER, S. J. *Wavelet Methods for Computer Graphics*. PhD thesis, Princeton University, January 1995.
- [14] HALSTEAD, M., KASS, M., AND DEROSE, T. Efficient, Fair Interpolation using Catmull-Clark Surfaces. In *Computer Graphics, Annual Conference Series, 1993* (1993), Siggraph, pp. 35–43.
- [15] JAFFARD, S., AND LAURENÇOT, P. Orthonormal Wavelets, Analysis of Operators, and Applications to Numerical Analysis. In *Wavelets: A Tutorial in Theory and Applications*, C. K. Chui, Ed. Academic Press, 1992, pp. 543–602.
- [16] LIU, Z., GORTLER, S. J., AND COHEN, M. F. Hierarchical Spacetime Control. In *Computer Graphics, Annual Conference Series, 1994* (August 1994), pp. 35–42.
- [17] LOUNSBERY, M., DEROSE, T., AND WARREN, J. Multiresolution Analysis for Surfaces of Arbitrary Topological Type. Tech. Rep. TR 93-10-05b, Department of Computer Science and Engineering, Princeton University, October 1993.
- [18] LYCHE, T., AND MORKEN, K. Spline Wavelets of Minimal Support. In *Numerical Methods in Approximation Theory*, D. Braess and L. L. Schumaker, Eds., vol. 9. Birkhauser Verlag, Basel, 1992, pp. 177–194.
- [19] MALLAT, S. G. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. *IEEE PAMI* 11 (July 1989), 674–693.
- [20] MEINGUET, J. Multivariate Interpolation at Arbitrary Points Made Simple. *Journal of Applied Mathematics and Physics (ZAMP)* 30 (1979), 292–304.
- [21] MORETON, H., AND SEQUIN, C. Functional Optimization for Fair Surface Design. *Computer Graphics* 26, 4 (July 1992), 167–176.
- [22] PENTLAND, A. Fast Solutions to Physical Equilibrium and Interpolation Problems. *The Visual Computer* 8, 5 (1992), 303–314.
- [23] QIAN, S., AND WEISS, J. Wavelets and the Numerical Solution of Partial Differential Equations. *Journal of Computational Physics* 106, 1 (May 1993), 155–175.
- [24] RANDO, T., AND ROULIER, J. Designing Faired Parametric Surfaces. *Computer Aided Design* 23, 7 (September 1991), 492–497.
- [25] SZELISKI, R. Fast Surface Interpolation Using Hierarchical Basis Functions. *IEEE PAMI* 12, 6 (June 1990), 513–439.
- [26] TERZOPOULOS, D. Image Analysis Using Multigrid Relaxation Methods. *IEEE PAMI* 8, 2 (March 1986), 129–139.
- [27] TERZOPOULOS, D. Regularization of Inverse Visual Problems Involving Discontinuities. *IEEE PAMI* 8, 4 (July 1986), 413–424.
- [28] WELCH, W., AND WITKIN, A. Variational Surface Modeling. *Computer Graphics* 26, 2 (July 1992), 157–166.
- [29] YSERENTANT, H. On the Multi-level Splitting of Finite Element Spaces. *Numerische Mathematik* 49 (1986), 379–412.