Compression Domain Rendering of Time-Resolved Volume Data

Rüdiger Westermann
Visualization and Media Systems Design
German National Research Center for Computer Science
Sankt Augustin, Germany

Abstract

An important challenge in the visualization of three dimensional volume data is the efficient processing and rendering of time-resolved sequences. Only the use of compression techniques, which allow the reconstruction of the original domain from the compressed one locally, makes it possible to evaluate these sequences in their entirety. In the following paper a new approach for the extraction and visualization of so called time-features from within time-resolved volume data will be presented. Based on the asymptotic decay of multiscale representations of spatially localized time evolutions of the data, singular points can be discriminated. Also the corresponding Lipschitz exponents, which describe the signals local regularity, can be determined, and can be taken as a measure of the variation in time. The compression ratio and the comprehension of the underlying signal will be improved, if we restore the extracted regions first, which contain the most important information.

Keywords and Phrases: *volume rendering, wavelet transforms, singularities, Lipschitz exponents*

1 Introduction

In recent years several methods have been developed to visualize static three dimensional volume data sets. Most of the proposed methods try to approximate more or less accurately the volume rendering integral [4][9][11]:

$$I(t_o, t_1) = \int_{t=t_0}^{t_1} q(t) e^{-\int_{s=t_0}^{t} \sigma(s)ds} dt \qquad (1)$$

where σ(s) defines the attenuation function, $q(t)$ is the volume source term, and t_0 and t_1 are the start and end points on the view ray. This reduced formulation of the more general and physically based transport equation [10], describes the summation of the light along a ray, which is scaled by a material dependent attenuation factor to get the final intensity. Basically, all methods can be classified into fast object space driven back-to-front projection methods [20][21], and in general slower but more accurate image space driven methods. The latter technique resamples the volume along the ray of sight, and is closely related to general integration rules, which evaluate the rendering integral numerically.

With the rapidly increasing resolution of the available data sets, it becomes more important to evaluate efficient compression techniques, which allow the computation of the rendering integral on the compressed domain. One possible solution is to evaluate the volume rendering integral on multiresolution representations of the original three dimensional signal [6][14][19]. Due to the sparse representation of projections into cascades of difference spaces, impressive compression ratios can be achieved. Furthermore, the rendering process can be performed on the compressed data directly, which avoids the memory consuming global reprojection of the data.

Even for the visualization of time-resolved volume sequences, the requirement to reduce the data as much as possible becomes most important. In this context, it is not only of special interest to process these time series in their entirety, but also to extract and enhance special features from within the sequences.

In the following, one new approach to deal with time-resolved volume data will be presented. The key idea is to compress each volume separately, but also to obtain information concerning the time evolution of certain parts of the volumes. This can be done by examining the time development of spatially localized multiresolution information. An additional compression of such regions that remain constant or vary slightly over time, and also an implicit description of the time dependence of certain regions is reached in this way.

The question that remains to be answered is how to extract the interesting information of the time varying volume sequences. This question implies that we define first the focus of interest of such sequences. In our approach we concentrate on the areas with strong variation over time. A measurement of this variation can be found by inspecting the signals wavelet transform. It has been shown that singular points within the data can be determined from the signals multiscale representation. Additionally, the Lipschitz exponents of such singularities can be measured from the asymptotic decay of the wavelet coefficients, as the scale goes to zero. The Lipschitz exponents characterize the local regularity of the signal, and can thus be used for the discrimination of those regions which should be reconstructed within high accuracy. In between these regions a less accurate reconstruction will be applied.

In the following, a short introduction to the theory of wavelet transforms is given, together with a short explanation of how to evaluate the volume rendering integral on multiresolution spaces. We will then focus on the singularity detection procedure and the integration into our approach will be outlined. Finally, some specific implementation details will be given, and results based on two data sequences will be compared.

2 Wavelet Transforms

Wavelet transforms can be seen as an effective tool to separate and analyse multiscale phenomena of the underlying data. Due to the fact that wavelets are well localized in both time and scale, they provide a useful method to examine these phenomena locally. This is the reason why wavelet transforms are of great interest in the field of volume rendering. Many researchers have investigated in more detail the basic concept and theory of wavelet transforms and multiresolution analysis over the past few years, and some introductions can be found in [8][1][12] and [17].

Basically, a wavelet decomposition is built up from scales and dilates of an infinite energy, self-similiar basis function $\Psi(x)$ with

$$\Psi_k^j(x) = \sqrt{2^j}\Psi(2^j x - k) \qquad \text{for } k \in Z.$$

At each scale 2^j, the translates of the scaled wavelet functions form a basis of some vector space W_j, which is the orthogonal complement of some approximation space V_j. These approximation spaces are built up from the scaling functions $\Phi_k^j(x)$, which generate a multiresolution analysis (MRA) of $L^2(R)$, if the nested sequence of subspaces $V_j = V_{j-1} \oplus W_{j-1}$ has certain properties.

In general, the so called difference spaces W_j contain the information that is lost when projecting a function from one approximation space V_j to the next coarser V_{j-1}. The projection operators for a certain function $f(x)$ into V_j and W_j respectively can be written as:

$$P^j f(x) = \sum_k S_k^j \Phi_k^j(x), Q^j f(x) = \sum_k D_k^j \Psi_k^j(x)$$

defining the smooth information $P^j f$ and the detail information $Q^j f$ which is needed to go from a certain resolution approximation space to the next finer one.

The wavelet coefficients S_k^j and D_k^j can be computed from the inner products [1] of the function with the dual scaling function and the wavelet:

$$S_k^j = \langle f(y), \tilde{\Phi}_k^j(y)\rangle, D_k^j = \langle f(y), \tilde{\Psi}_k^j(y)\rangle$$

An efficient method to perform the inner product calculations has been proposed by Mallat, whose pyramid algorithm runs in linear time with the number of function samples [12]. Basically, this algorithm stems from the so called *two-scale* relation, defining the basis functions on a certain level as linear combinations of basis functions on the next finer one. This combination can be expressed with some *low pass* and *high pass* filter sequences, which have do be applied recursively on the smooth approximations of the signal.

For the case of separable MRA's the extension to higher dimensions can be easily done, constructing the three dimensional basis functions from the tensor products of the one dimensional ones.

An important property of the wavelet transform is, that the original signal can be reconstructed from its expansion coefficients locally:

$$f(x) = \sum_{j,k} D_k^j \Psi_k^j(x)$$

Taking this reconstruction property into account, the volume rendering integral can be evaluated on the projected signal directly. The major advantage of this change of basis stems from the sparseness of the projections into the difference spaces. Based on the number of vanishing moments of the wavelet, which is the largest number p for which all integrals

$$\langle x^n \rangle_\Psi = \int_{-\infty}^{\infty} \Psi(x) x^n \, dx \qquad \text{n} = 0..p-1$$

vanish, polynomials up to a degree of $p-1$ can be reproduced without error within the approximation spaces only. This implies that a large number of coefficients will be zero or less than a predefined threshold, and can be neglected without increasing a given error tolerance.

3 Multiscale Singularity Detection

For the characterization of signals it is often of major interest to discriminate singular points of the signal, and to determine those parts where the signal behaves in a less regular manner. This is due to the fact that singularities or irregular structures contain most of the interesting signal information, and often allow one to determine the important features.

[1] An inner product of two functions is defined as:

$$\langle f(u), g(u)\rangle = \int_{-\infty}^{+\infty} f(u) g(u) du$$

Additional information is given, if one is able to measure the regularity of the signal at these singularities. This can be done with Lipschitz exponents, which define the local regularity of functions.

Given a subset $S \subset D(f)$ and a constant $M > 0$, a function $f(x)$ is said to be Lipschitz $\alpha \geq 0$ over S, if for two points $t, \tilde{t} \in S$ the following holds:

$$|f(t) - f(\tilde{t})| \leq M|t - \tilde{t}|^{\alpha}$$

For this case we write $f \in LIP_S(M, \alpha)$.

As a result it can be shown that $f(x) \in C^0$ over S, if $f \in LIP_S(M, \alpha)$. If $\alpha > 1$ over a certain interval $[a, b]$, and $f \in LIP_{[a,b]}(M, \alpha)$, then f is constant over $[a, b]$. If the derivative of $f(x)$ is bounded by some constant K, that is $|f'(t)| \leq K$, and $f \in C^1$ over an interval $[a, b]$, then $f \in LIP_{[a,b]}(M, 1)$.

In general, increasing Lipschitz exponents indicate a more regular behaviour, and there is also a strong relationship between the differentiability of $f(x)$ and the Lipschitz regularity. If $f(x) \in LIP_S(M, \alpha)$ at a point x_0 with $\alpha > n$, then f is n times differentiable at x_0. For functions that are discontinous but bounded at a certain point, the Lipschitz exponent equals 0 at this point.

Another interesting case occurs at sharp peaks, where the signal tends to be more singular than discontinous. In this case the Lipschitz regularity is negative at this point.

The extraction of the signals singularities and the computation of corresponding Lipschitz exponents, permits the classification of parts of the signal, and the distinction between these parts by considering their local regularity. The question that remains to be answered is how to detect the singularities and how to compute the Lipschitz exponents.

Jaffarth [7] proved a general theorem, which states that the singularities of a signal can be detected from its multiscale representation, and that the local Lipschitz regularity can be computed from the decay of the wavelet transform across scales.

The key idea is to measure the asymptotic behaviour of the wavelet coefficients as the scale goes to the finest resolution. Formally, the relation between the absolute values of the wavelet transform and the Lipschitz regularity at a point x_0 can be described as:

$$|Wf(s, x)| \leq A * (s^{\alpha} + |x - x_0|^{\alpha}) \tag{2}$$

where $Wf(s, x)$ describes the wavelet transform on a certain scale s. This relation holds for points x within a neighborhood of x_0, but depending on the location of x, the decay of the multiscale coefficients behaves differently. For points x at a certain scale s, with $|x - x_0| \leq s$, (2) implies a $O(s^{\alpha})$ decay, whereas for other points the decay is controlled by their distance to x_0. For practical computations of the Lipschitz exponents, Mallat et al. [13] proposed an efficient method, which is based on the observation, that the local maxima of the wavelet transform on every scale define *scale-space curves*. Connecting those maxima which proceed from a certain scale to the next coarser one, the resulting maxima lines point exactly to the singular points of the signal on the finest scale (see Figure 1). For isolated singularities the local Lipschitz regularity at a point x_0 can now be determined from the asymptotic decay of the coefficients along the maxima lines pointing to x_0. For a function to be $\in LIP_S(M, \alpha)$ at x_0, it has to be hold that the decay behaves as:

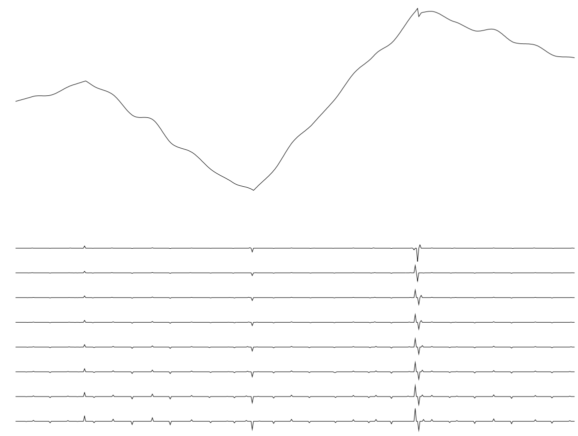

Figure 1: Original signal and local maxima of the wavelet transform. One clearly realizes the maxima lines pointing to the singular points.

$$|Wf(s, x)| \leq K * (s^{\alpha}) \tag{3}$$

K and α can then be measured, only taking into consideration the local maxima along the scale-space lines. This allows the discrimination of singular points, and the separation of parts of the signal taking into account its local behaviour and characteristics.

4 Time-Feature Extraction

Once we are able to detect the singularities of a given signal as well as to characterize their regularity, the question is how to integrate this method into a multiresolution approach for the visualization of time-resolved volume data. Since we wish to extract those features from the volume sequence that vary strongly over time, the key idea is to separate spatially corresponding regions from within each volume, and to examine their relation in more detail.

Instead of applying the mentioned approach to the time evolution of each voxel separately, we relate the time evolution of corresponding multiscale components, respectively wavelet coefficients (see Figure 2).

Basically, this approach has two main advantages:

- In a pre-processing step each volume can be separately transformed, and compressed with respect to its sparse multiresolution representation. During this computation, no information from neighboring volumes is needed, and an optimal compression ratio for each time step is reached. In this way we avoid the storage of the

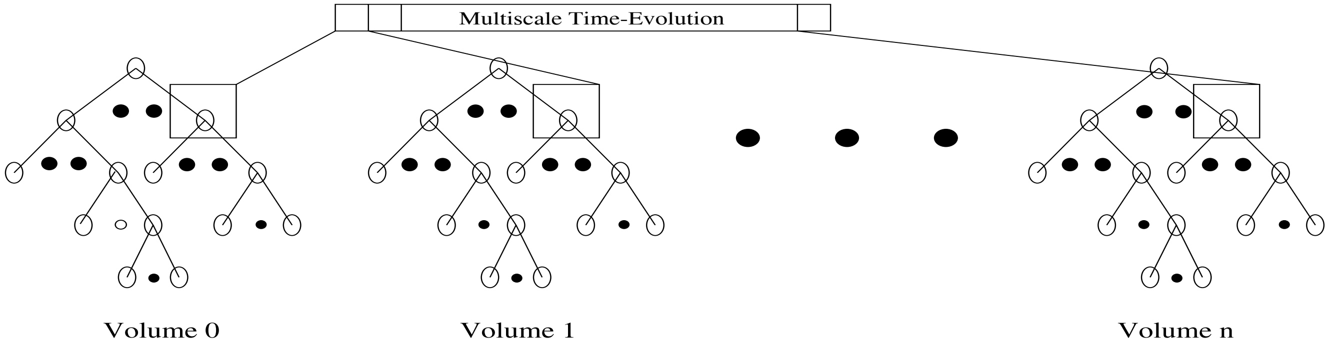

Figure 2: Tree like data structure to store multiresolution information for each time step. Corresponding multiresolution information is reorganized as a one dimensional signal, to obtain the time evolution of certain regions.

whole sequence, and we always minimize the memory requirement. Furthermore, a large number of entries for each wavelet coefficient time evolution will be zero. This is due to the sparse representations of changes in time, even for scarcely varying structures.

- Due to the fact that we examine the relation between multiscale components of each volume, we divide low frequency regions, which remain constant over time within each time step, from those high fluctuating parts, which vary strongly. This allows us to separate turbulent or irregular structures as well as to determine those structures, which are present within all volumes. This regions can be found by considering all coefficients which remain constant within all multiresolution representations.

Once the compressed multiscale structures for each volume have been constructed, we examine the time evolution of each component successively, starting at the finest level. For each of the generated one dimensional signals, which contain information concerning the time behaviour of a certain region, an additional wavelet expansion has to be computed. The singularity detection process can now be applied on the generated multiscale representation directly. Inequality (3) implies that one has to find those K and α, which lead as close as possible to an equality for each scale s. These values can be found by minimizing the derived "energy"-function

$$\sum_{s=1}^{s_{max}} (log(|Wf(s,x)|) - log(K) - \alpha * log(s))^2 \qquad (4)$$

Both, the function value itself and also its first derivative, can be calculated for arbitrary K and α, so that a steepest descent algorithm to solve the minimization problem can be used. Starting at a random point on the finest scale, successive steps into the direction of the steepest gradient have to be performed to find the next point. Of course this method will not necessarily find the global minima, but it is often sufficient only to determine the overall behaviour of the decay along the maxima lines. In practice, this computation has only to be done for a small number of points on the finest scale, and some kinds of coarse pre-separation allow one to distinguish between specific classes of singularities. In general, increasing coefficients indicate sharp singular peaks with Lipschitz exponent less than zero, discontinous regions are classified by coefficients that remain constant across scales which implies Lipschitz exponent equal to zero, while more regular behaviour is indicated by increasing Lipschitz exponents. Furthermore, one can assume that for sharply varying regions, the coefficients on all scales have to exist. For example, Mallat et al. [13] proved that random noise can be separated from the original signal, extracting those singularities for which the scale-space maxima lines disappear at the next scale.

Considering the evolution of the wavelet coefficients across scales, we are able to discriminate points where the time course behaves singular, and also to estimate the Lipschitz regularity from the asymptotic decay as the scale goes to the finest resolution. These computed values indicate those singularities which are of interest and define the important time-features within the volume sequence.

In our approach we take the singular points for which the Lipschitz regularity is not greater than a given bound, to determine which regions to reconstruct most accurately. In between these regions we increase the error tolerance for the reconstruction process, which allows us to eliminate a larger number of coefficients. This method is closely related to the wavelet probing algorithm, proposed by Deng et al. [3]. They split up the signal into smooth segments that can be compressed separately, based on wavelets which are defined over an interval. Compared to our approach, their method also prevents discontinuities to be blurred, but instead of separating smooth segments in between the discontinuities, we prefer to separate the discontinuities themselves.

Another approach is to enhance only those features which behave irregularly over a certain time interval, while other regions will be completely neglected. Based on the generated multiscale representation, this method implies the storage of those coefficients, which are needed to reconstruct the irregular structures without error. All other parts will be neglected. Due to the increasing support of the basis functions on the coarser level, they will overlap most of the singular points, and the smooth approximations of the signal in between these points are retained.

In general, the additional transformation and singularity processing of the time evolutions leads to increased compression of the time domain. This allows the processing of large scale time-resolved sequences in their entirety on standard workstations without additional memory.

5 Implementation

One major goal of our implementation is to integrate the earlier proposed multiresolution framework [19] into the current approach for the visualization of time-resolved volume sequences. To visualize volume data which corresponds to a certain time step, the multiresolution representation for that time step has to be reconstructed from the transformed time sequence first. Each wavelet coefficient for that time step has to be reconstructed from the already compressed multiscale time evolution of this coefficient. The processing of these time evolutions, including the singularity detection and the determination of the Lipschitz regularity, is performed for each coefficient separately in a pre-processing step. As a result, we obtain a multiresolution representation for each coefficient, which describes the fluctuations over time at different scales. The standard reduction and compression techniques (thresholding, quantization, coding) can then be applied to each multiscale fluctuation separately.

This approach increases the numerical complexity of the rendering process to some extent, but due to the interpolating properties of the basis functions used, arbitrary interpolations between successive data sets can be reconstructed. Even for static regions the reconstruction process will be accelerated, due to the sparseness of the multiresolution representation.

5.1 Choice of basis

Two basic requirements have to be taken into account in choosing the optimal basis functions for our implementation. On the one hand, the support of the basis functions used should be as compact as possible, accelerating the reconstruction process as much as possible. On the other hand, the number of vanishing moments of the wavelets used should be as high as possible, leading to better compression ratios. For the expansion along the time axis, we are also interested in good interpolation properties of the basis functions. For the generation of smooth interpolants between successive time steps, the basis functions should have good regularity properties and should oscillate as little as possible between given samples.

Another important feature that has to be considered, is the maximal Lipschitz regularity we want to be able to determine. Basically, one needs to use wavelet basis functions with more than α vanishing moments to characterize singularities of Lipschitz regularity α. Only in this case the Lipschitz regularity of the function can be measured from the decay of the wavelet coefficients along the maxima lines.

Based on these preconditions, we decided to use different wavelets for the construction of the three dimensional multiresolution representations of each static volume, and for the additional processing of the time evolutions of each multiscale component. For the first task we used Daubechies wavelets with two vanishing moments (see Figure 3). They are the most compact orthogonal ones with respect to the number of vanishing moments. Due to the orthogonality of the Daubechies wavelets, they are dual to each other, and the same basis functions can be used for the construction and reconstruction process.

For the singularity detection process we used non-orthogonal, but compactly supported B-Spline wavelets. This choice was motivated by the regularity property of polynomial splines, which are indeed the interpolants which oscillate the least among all other interpolants for a certain degree. Unser et al. [18] proposed a general framework for the generation of polynomial splines of certain degree and the corresponding filter sequences for the wavelet expansion. This allows us to increase the number of vanishing moments arbitrarily. Due to the higher Lipschitz regularity that can be characterized, polynomial splines can be adapted efficiently to the kinds of irregular structures one wants to be able to detect.

The disadvantage of standard B-Splines is that they are not orthogonal, and that the corresponding dual basis functions are of infinite support. This leads to infinite filter sequences, either for the construction, or the reconstruction process. As we are mostly interested in simplifying the reconstruction process as much as possible, it has to be done with the compactly supported B-Splines wavelets. To avoid the infinite filter sequences for each stage of the construction process, we first project the original signal in the dual basis, which allows to use the compactly supported filter sequences for the wavelet transformation of the one dimensional time evolutions. Finally, we perform an additional change of basis of the multiresolution information, reprojecting into the standard B-Spline basis. Because the duals span the same residual spaces as the standard B-Splines do, we do not have to worry about errors during the projections from one basis to another. Care has to be taken for the accurate evaluation of the projection operators. As the corresponding filter sequences are infinite, boundary problems and the error due to the necessary cut off of the sequences have to be taken into consideration. We used an inverse mirroring at the start and end points of the signal, and filter entries below a certain threshold are neglected.

Within our implementation cubic B-Spline wavelets (see Figure 4) with 2 vanishing moments have been applied for the examination of the time developments of the multiscale components. As a result, sharp variation points such as peaks, discontinuities or turbulent structures with a Lipschitz regularity between 0 and 2 can be detected. A good localization of varying structures has been reached in this way.

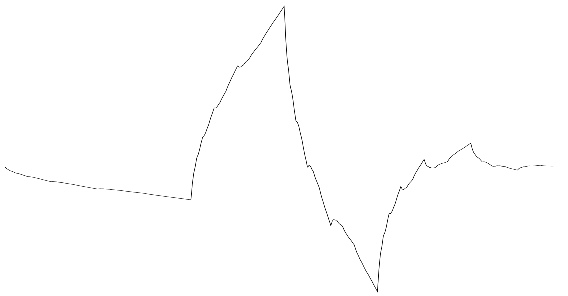

Figure 3: Daubechies wavelet with 2 vanishing moments. The support of the basis function is 3.

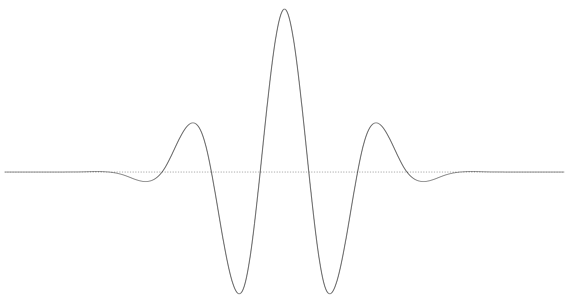

Figure 4: Compactly supported cubic B-Spline wavelet. The number of vanishing moments is 2 and the support is 6.

6 Results

The implementation has been tested on two volume sequences. The first data sequence, which was built up from 64 128^3 volume data sets, was generated from a time varying stochastic fractal, simulating three dimensional cloud structures. The generation process is based on Perlins hypertexture function [16], adding multiple scaled and dilated copies of a random noise function, to obtain the typical $1/f^\beta$ spectrum of stochastical fractals. Based on Taylors frozen turbulence hypothesis [5], the variations of the fractal in time are equivalent to its spatial behaviours and properties. This observation allows the generation of each time step as a snapshot of a four dimensional stochastic fractal. To show as well as possible the application of the proposed time-feature extraction technique, we modified the generation process locally. Away from the center of the volumes, we decreased the number of summed octaves used for the simulation of the fractal in time, neglecting more and more high frequencies. Within a sphere around the midpoints we simulated the time behaviour on full resolution, which leads to high fluctuating structures inside the sphere, whereas outside the sphere all structures remain constant. In this way, we obtain for each time step a three dimensional fractal with equivalent properties within the whole volume, but with different properties for certain regions in time. Figure 5 shows a snapshot of an arbitrary time step from this sequence. The top left side shows the reconstructed volume, rendered on full resolution out of the whole sequence with 512 MB. On the top right side the same time step is reconstructed out of the already compressed sequence, which consists only of 74 MB. One can clearly observe that almost all information is retained, whereby for this specific example the compression of the time evolutions saved 150 MB of memory. The rendering method we used for this example is well adapted to the visualization of low albedo, low density volumes, summing up the density values along the ray only. Figure 5 below shows the same time step as above, but visualized based on the described time-feature extraction method. On the bottom left side we retain all coefficients, which contribute to the detected singularities in time, up to a Lipschitz regularity of 2. Due to the fact that on the coarser resolution level the static regions are influenced by coefficients that belong to the turbulent areas, the smooth approximations of the volumes contribute to these static regions also. It can be realized away from the center, where the structures are more and more blurred. This behaviour leads to better compression ratios, and also allows one to discriminate those structures, which vary strongly over time. The application of the time-feature enhancement approach is shown on the bottom right side of Figure 5. We only reconstructed those structures accurately, which behave irregular over time. All other regions, for which we do not determine any singularities within their time evolution, are suppressed. The time varying structures are retained completely within the sphere around the center, whereas outside the sphere all information is neglected.

The second time series was built up from 128 64^3 volumes, which are homogeneous everywhere, but a high density object, moving along a sine wave from one corner of the volume to the opposite. An interesting behaviour of the multiscale processing of each time evolution of certain regions appears, if we use only the coarsest levels for the reconstruction of each time step. In our example we took only the 3 (left side of Figure 6), respectively 2 (right side of Figure 6) coarsest levels out of 7. Due to the low pass filtering in the time domain, more and more global information from neighboring time steps is projected into each volume. The amount of information that is integrated from other time steps is controlled by the level up to which the reconstruction is performed. Using the coefficients from all resolution levels, each volume can be rebuilt completely. With decreasing resolution level, more and more neighbors contribute to the final volume. This can be seen on the right of Figure 6, where we obtain information from the whole sequence within one static volume. Basically, this technique can be seen as a blurring in the time domain, which allows to examine the location of certain objects at different time steps and their movement over time in more detail.

7 Conclusion

A memory minimizing approach for the visualization of time-resolved volume data based on multiresolution representations of both the spatial distribution and the time evolutions of multiscale components has been proposed. With respect to the evolution of certain features across scales, those regions that vary strongly over time can be separated and rendered prominently, whereas scarcely varying structures can be suppressed. This leads to an additional compression of the overall amount of data, and enhances the understanding of a wide variety of special kinds of volume data to some extent. Even for the extraction of irregular structures from time varying fluid fields, but also for the discrimination of static parts within time-resolved sequences, the presented method leads to good results. Of special interest within this context is the extraction and prediction of motion paths for certain well defined objects. Due to the fact that the wavelet transform splits up the frequency domain into dyadic frequency bands, a careful examination and comparison of the time behaviour of corresponding bands could give some new insights to these kinds of problems.

The main disadvantage of the proposed method is the increasing numerical complexity of the rendering process. In addition to the time consuming evaluation of the rendering integral based on the wavelet expansion of the original volume, a further reconstruction step has to be done, which generates for each time step the multiresolution representation. Additionally, a complex pre-processing step has to be performed, determining the singularities and the local Lipschitz regularity. On the other hand this process has to be done only once, and due to the spatial localization of the wavelet transform, the presented approach can be easily implemented on distributed memory, parallel computers. Dividing the volume into subblocks, which will be distributed over the nodes, the wavelet transform of each of these blocks can be performed individually. Of course this forces some global communication depending on the length of the filter sequences, but once the multiresolution representation for a certain block has been computed, equivalent subblocks for the next time steps are stored on the same node. The singularity detection procedure based on the time evolutions of the subblocks wavelet coefficients, as well as the final compression of these blocks, can then be done locally.

8 Acknowledgement

The author wish to thank Wim Sweldens and Heinrich Müller for valuable discussions and advises.

REFERENCES

[1] C. K. Chui. Wavelets-A Tutorial in Theory and Applications. Academic Press, pages 91-122, 1992.

[2] I. Daubechies. Ten Lectures on Wavelets. In *CBMS-NSF Series in Applied Mathematics*, Vol. 61, SIAM 1992.

[3] B. Deng, B. Jawerth, G. Peters and W. Sweldens. Wavelet Probing for Compression based Segmentation. In *SPIE Conference Proceedings*, 1993.

[4] R. Drebin, L. Carpenter and P. Hanrahan. Volume Rendering. In *Computer Graphics*, Vol. 22, No. 4, pp. 65-74, 1988.

[5] W. Frost and T. Moulden. Handbook of Turbulence, Vol. 1: Fundamentals and Applications. Plenum Press, New-York, 1977.

[6] M. Gross, L. Lippert, A. Dreger and R. Koch. A new Method to Approximate the Volume Rendering Equation Using Wavelets and Piecewise Polynomials. In *Computers and Graphics*, Vol. 19. No. 1, 1995.

[7] S. Jaffarth. Exposants de Holder en des points donnes et coefficients d'onolettes. In *Notes au Compte-Rendu de l'Academie Des Sciences*, No. 308, I, pages 79-81, 1989.

[8] B. Jawerth and W. Sweldens. An Overview of Wavelet Based Multiresolution Analysis. In *SIAM Review*, Vol.36, No.3, 1994.

[9] J. Kajiya and B. Herzen. Ray tracing volume densities. In *Computer Graphics Proceedings, Annual Conference Series*, pages 165-174, 1984.

[10] W. Krueger. The Application of Transport Theory To Visualization of 3D Scalar Data Fields. *IEEE Visualization'90, Conference Proceedings*, pages 12–15, 1990.

[11] M. Levoy. Display of Surface from Volume Data. In *IEEE Computer Graphics and Applications*, Vol.8 No.3, pages 29-37, 1988.

[12] S. Mallat. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence* Vol.11 No.6, pages 674-693, 1989.

[13] S. Mallat and W. Hwang. Singularity Detection and Processing with Wavelets. In *IEEE Transactions on Information Theory*, Vol.38, No.2, pages 617-643, 1992.

[14] S. Muraki. Approximation and Rendering of Volume Data Using Wavelet Transforms. In *IEEE Computer Graphics and Applications*, Vol.13, No.4, pages 50-56, 1993.

[15] S. Muraki. Multiscale 3D Edge Representation of Volume Data by a DOG Wavelet. In *Proceedings ACM Workshop on Volume Visualization*, pages 35-42, 1994.

[16] K. Perlin and E. Hoffert. Hypertexture. In *Computer Graphics Proceedings, Annual Conference Series*, pages 253-263, 1989.

[17] G. Strang. Wavelets and Dilation Equations : A brief Introduction. In *SIAM Review*, Vol.31, No.4, pages 614-627, 1989.

[18] M. Unser and A. Aldroubi. Polynomial Splines and Wavelets-A Signal Processing Perspective In *Wavelets-A Tutorial in Theory and Practice*, Academic Press, pages 91-122, 1992.

[19] R. Westermann. A Multiresolution Framework for Volume Rendering. In *Proceedings ACM Workshop on Volume Visualization*, pages 51-58, 1994.

[20] L. Westover. Footprint Evaluation for Volume Rendering, In *Computer Graphics Proceedings, Annual Conference Series*, pages 185–188, 1990.

[21] J. Wilhelms and A. van Gelder. A Coherent Projection Approach for Direct Volume Rendering. In *ACM Computer Graphics*, Vol.25, No.4, 1991.

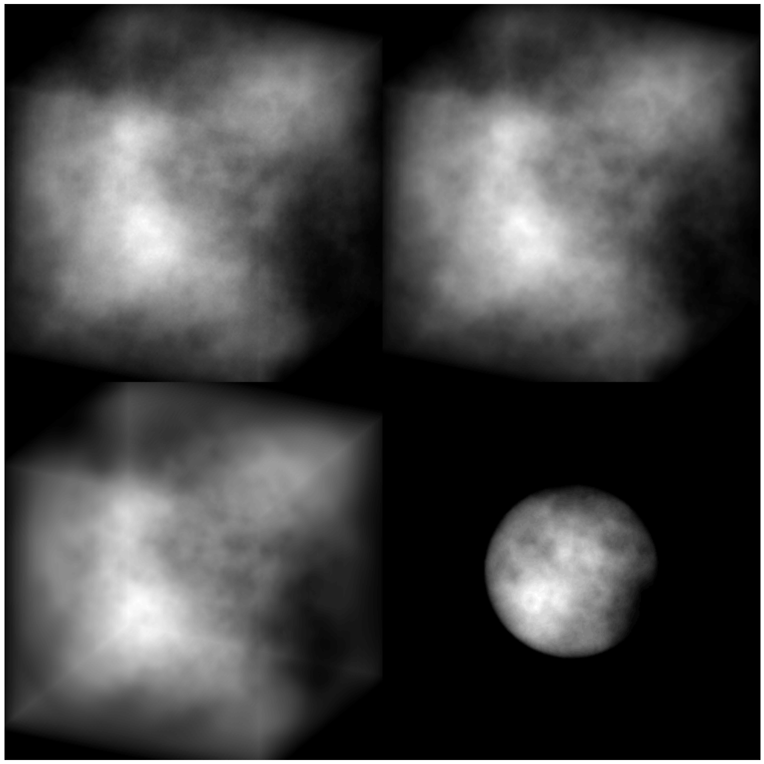

Figure 5: Stochastic Fractal: top left shows a time step rendered on full resolution from the sequence with 512 MB; top right shows the same time step rendered from the compressed sequence with 74 MB; bottom left shows the time step reconstructed only from those coefficients, which contribute to varying structures; bottom right shows the reconstruction of the varying structures only.

Figure 6: Moving object: left shows the reconstruction of a certain time step from the 3 coarsest resolution level; right shows the same time step reconstructed from the 2 coarsest resolution level.