Published in final edited form as:

Comput Vis Image Underst. 2004 December 1; 96(3): 435–452. doi:10.1016/j.cviu.2004.03.010.

Volumetric Video Compression for Interactive Playback★

Bong-Soo Sohn, Chandrajit Bajaj, and Vinay Siddavanahalli

Department of Computer Sciences and ICES, University of Texas, Austin, TX 78712

Bong-Soo Sohn: bongbong@cs.utexas.edu; Chandrajit Bajaj: bajaj@cs.utexas.edu; Vinay Siddavanahalli: skvinay@cs.utexas.edu

Abstract

We develop a volumetric video system which supports interactive browsing of compressed time-varying volumetric features (significant isosurfaces and interval volumes). Since the size of even one volumetric frame in a time-varying 3D data set is very large, transmission and on-line reconstruction are the main bottlenecks for interactive remote visualization of time-varying volume and surface data. We describe a compression scheme for encoding time-varying volumetric features in a unified way, which allows for on-line reconstruction and rendering. To increase the run-time decompression speed and compression ratio, we decompose the volume into small blocks and encode only the significant blocks that contribute to the isosurfaces and interval volumes. The results show that our compression scheme achieves high compression ratio with fast reconstruction, which is effective for client-side rendering of time-varying volumetric features.

Keywords

Time-Varying Volume Visualization; 3D Video; Compression; Isocontouring; Hardware-Acceleration; Wavelet Transform

1 Introduction

Scientific simulations of today are increasingly generating densely sampled time-varying volume data which have very large sizes. For example, the size of an oceanographic temperature change data set tested in this paper is 237MB/frame ($2160 \times 960 \times 30$ float) $\times$ 115 frames, and the gas dynamics data set is 64MB/frame ($256 \times 256 \times 256$ float) $\times$ 144 frames. To visualize such time-varying volumetric data, *volume rendering* and *isocontouring* techniques are performed frame by frame, so that a user can navigate and explore the data set in space and time. Both rendering techniques have their own strengths. While volume rendering can display amorphous volumetric regions specified by the transfer function with transparency, isocontouring can provide the geometric shape of the surfaces specified by significant isovalues. Both techniques can be combined for better understanding of the overall information present in a volumetric data set. For example, figure 1 shows visualization of simulated explosion during galaxy formation through rendering of time-varying volumes and isosurfaces.

While current state-of-the-art graphics hardware allows very fast volume and surface rendering, transfer of such large data between data servers and browsing clients can become a bottleneck due to the limited bandwidth of networks. Efficient data management is an important factor in the rendering performance. To reduce the size of the data set, it is natural to exploit temporal and spatial coherence in any compression scheme. However, since the data size of

★Visit <http://www.ices.utexas.edu/~bongbong/volvideo> for accessing video files.

even a single frame is very large, run-time decompression can also be a bottleneck for interactive playback.

From this motivation, we have developed a unified compression scheme for encoding time-varying volumetric features. In most cases, we use the term *feature* to mean an isosurface and/or an interval volume specified by a scalar value range. An interval volume can be a whole volume. The inputs for compression are k iso-values iso_i , $i = 1..k$, l value ranges $[r_a, r_b]_i$, $i = 1..l$ and discretized time-varying volume data V . The data V , containing T time steps can be represented as $V = \{V_1, V_2, \dots, V_T\}$, where $V_t = \{f_{i,j,k}^t | i, j, k \text{ are indices of } x, y, z \text{ coordinates}\}$ is the volume at time step t , containing data values $f_{i,j,k}^t$ at the indices i, j, k . Our primary goal in compression is to

- compress time-varying isosurfaces and interval volumes in a unified way.
- reduce the size of time-varying volumetric data with minimal image degradation.
- allow real-time reconstruction and rendering with PC graphics hardware acceleration.

We borrow the idea of MPEG compression to efficiently exploit spatial and temporal coherence in data sets. However, direct extension of MPEG for 2D video compression to time-varying isosurface and feature compression is not suitable for satisfying our goals. Since MPEG encodes and decodes every block in each frame of the volumetric time-varying data, unnecessary regions within the data may also be encoded and decoded.

We adopt a block-based wavelet transform with temporal encoding in our compression scheme. The wavelet transform is widely used for 2D and 3D image compression. By truncating insignificant coefficients after wavelet transformation, these schemes achieve high compression ratio while keeping minimal image distortion. However, complete transformation of each frame is a waste of space and time resources, because function values not contributing to the given isosurfaces and volumetric features do not need to be encoded. In addition, the contributing values which have small changes over time do not need to be updated. Therefore encoding and decoding only the values significant in space and time instead of the full volume can improve both compression ratio and decompression speed.

Each volumetric frame is classified as either an intra-coded frame or a predictive frame. The intra-coded frames can be decompressed independently while the predictive frames are the differences from their previous frames. Assuming that different blocks have different temporal variance, we can sort the blocks based on their temporal variance and truncate insignificant blocks to achieve higher compression ratios and faster decompression speeds.

In addition to efficient compression, fast reconstruction and rendering of the iso-surface and volumetric features are also achieved. By attaching seed cells in the compressed stream of each volume frame, the rendering browser can construct the isosurfaces in minimal time. The fast speed comes from removing the search phase for finding at least one cell intersecting with each isosurface component. Since the isosurface is fixed, we need to store only one seed cell per isosurface component, which hardly affects the compression ratio. We can also compress the selected set of components in isosurfaces and volumes, and their evolution by using the feature tracking method [24]. The reconstructed features can be rendered in real-time using PC graphics hardware.

Figure 2 shows the overall architecture of our interactive volumetric video display system. The algorithm requires the features to be defined before compression. These features can be identified manually or by automatic feature detection tools [18]. To save disk storage space and to overcome the limitation of I/O bandwidth in network systems, a series of compressed

frames are read from data source servers to browsing clients. Once each compressed frame is read, it is decompressed in software. The reconstructed image array is used for accelerated isocontouring and also sent to the texture memory in the graphics hardware for displaying volumetric features. This architecture can allow users to explore and interact with iso-surfaces embedded in the amorphous volumetric features in space and time, which is our ultimate goal.

The remainder of this paper is organized as follows. First, the related work is described. Then, in section 3, an architecture for displaying volumetric video is described. In section 4, a volumetric video compression scheme supporting interactive decompression is proposed. Section 5 describes our scheme for interactive browsing of compressed time-varying features. Experimental results are described in section 6. Finally, in section 7, we give a conclusion.

2 Related Work

Visualization of time-varying volume data has been a challenging problem. Compression, time-based data structures, and high performance visualization systems have been introduced to cope with overwhelming data sizes and heavy computation requirements.

Compression

Compression is extremely useful for large data manipulation, especially for transmission of data from servers to browsing clients. Since scientific data tend to be very large and have lots of redundancy, people prefer to use compressed data for efficient use of the memory and I/O bandwidth. A number of algorithms for image and surface compression have been developed.

Papers on single resolution and progressive compression of triangulated surfaces (e.g. isosurfaces) include those by [27,12]. A compression scheme specialized for isosurfaces [29] utilizes the unique property of an isosurface that only the significant edges and function values defined on a vertex are required to be encoded.

Most image compression techniques are geared towards achieving the best compression ratio with minimal distortion in the reconstructed images. JPEG and MPEG [6] are developed for compressing still images and 2D video data with controllable size and distortion tradeoff. Embedded coding algorithms such as embedded zero tree wavelet (EZW) [21] and set partitioning in hierarchical trees (SPIHT) [19] are useful for progressive transmission and multimedia applications.

For 3D image compression, Ihm et al. [11] described a wavelet-based 3D compression scheme for visible human data and later extended it to 3D RGB image compression for interactive applications such as light-field rendering and 3D texture mapping [2]. Compression ratios can be improved by capturing and encoding only significant structures and features in the data set [1,16]. In those compression schemes the primary goal is fast random access to data, while maintaining high compression ratios. This allows interactive rendering of large volume data sets [9].

Guthe and Straser [8] applied the MPEG algorithm to time-varying volume data using wavelet transformations. They compare the effects of motion compensation and the usage of different wavelet basis functions. Lum et al. [14] exploits texture hardware for both rendering and decompression. Since data is transferred in the compressed format between different memory systems, I/O time is significantly reduced.

Time Based Data Structures

Time-Space Partitioning(TSP) Tree was introduced and accelerated later by using 3D texture mapping hardware [5] for fast volume rendering of time-varying fields. The efficiency comes

from skipping insignificant rendering operations and reusing the rendered images of the previous time step.

Shen[22] proposed the Temporal Hierarchical Index (THI) Tree data structure for single resolution isocontouring of time-varying data, by an extension of his ISSUE algorithm [23]. The THI tree provides a compact search structure, while retaining optimal search time. Hence, expensive disk operations for retrieving search structures are reduced. Sutton et al. [26] proposed Temporal Branch-on-Need Tree by extending octrees for minimizing unnecessary I/O access and supporting out-of-core isosurface extraction in time-varying fields.

Shamir et al. [20] developed an adaptive multi-resolution data structure for time dependent polygonal meshes called *T-DAG* (Time-Direct Acyclic Graph). *T-DAG* is a compact representation which supports queries of the form (*time-step*, *error-tol*), and returns an approximated mesh for that time step, satisfying the error tolerance.

High Performance Visualization Systems

Ma and Camp [15] describes a remote visualization system under the wide area network environment for visualization of time-varying data sets. Current state-of-the-art graphics hardware enables real-time volume and isosurface rendering [28] and decompression [14].

Isosurface Extraction

A large amount of research has been devoted in the past for fast isosurface extraction from 3D static volume data. The Marching Cubes algorithm [13] visits each cell in a volume and performs appropriate local triangulation for generating the isosurface. To avoid visiting unnecessary cells, accelerated algorithms [4] minimizing the time to search for contributing cells are developed.

The contour propagation algorithm is used for efficient isosurface extraction [3,10]. Given an initial cell that contains an isosurface component, the remainder of the component can be traced by contour propagation. This property significantly reduces the space and time required for searching cells containing the isosurface by using a small number of seed cells. Multiresolution [7] and view-dependent techniques [30] are useful to reduce the number of triangles in an isosurface.

3 Interactive Volumetric Video

Like widely used 2D video systems, a *volumetric video* system displays a sequence of 3D images over time, frame by frame. While in a 2D video, users can only look at continually updated 2D images in a passive way, a volumetric video, or time-varying volume visualization system allows them to explore and navigate the 3D data in both space and time. Considering that most scientific simulations generate dynamic volume data, volumetric video systems are especially helpful for scientific data analysis.

A naive way for displaying time-varying 3D volume data is to repeat reading each frame from the data server and rendering the volume with the given visualization parameters. Since most time-varying scientific data sets are very large and have high spatial and temporal coherence, it is natural to apply compression for reducing storage overheads and transmission times. However, run-time decompression of data encoded by standard static and time-varying image compression schemes may become a bottleneck in real-time playback of volumetric video because they usually decompose an image into blocks and decode every block during decompression. During compression, we order the blocks based on their significance and encode only significantly changing blocks. This increases the run-time decompression speed as we have limited the number of blocks to decode.

A two-stage strategy is adopted to enable interactive navigation and exploration of very large time dependent volume data. In the first stage on the server side, the large time-dependent volume is analyzed and processed on a high performance server so that results of this volumetric processing is an intermediate multi-resolution, time-dependent volumetric representation of interesting features (isosurface, volume within a range) of the data, generated and stored in a compressed format. The intermediate multi-resolution representation permits tradeoffs between interactivity and visual fidelity for the second, interactive browsing stage. In the second stage on the client side, the volumetric video is decoded and played back by an interactive visualization browser that can be made available on a standard desktop workstation equipped with a 3D graphics card. In contrast to a standard video player, the visualization browser can allow certain levels of interactivity such as dynamically changing viewing parameters, modifying lighting conditions, and adjusting color-opacity transfer functions, in addition to timed playback of the volumetric video along with some user-specified fly-path in space and time.

4 Compression Scheme

In this section, we describe a unified scheme for compressing both time-varying isosurfaces and volumetric features at the same time. By encoding only significant function values based on associated weights using a wavelet transform, we can achieve high compression ratios. However, simple function encoding requires online isosurface extraction in the client side. To accelerate this surface extraction process, we insert seed cells into the compressed volume frames.

4.1 Compression

The input data set is a time dependent volume data set, $V = \{V_1, V_2, \dots, V_T\}$ with k isovalues $iso_1, \dots, iso_k$ and l value ranges $r_1 = [r_a, r_b]_1, \dots, r_l = [r_a, r_b]_l$. Each frame of the volume is classified as either an intra-coded frame or a predictive frame. For each isovalue and range, the reconstructed quality can be specified as a threshold value for wavelet coefficients and another threshold value for the blocks in predictive frames, such that wavelet coefficients or blocks not satisfying that value are truncated. The criteria for truncation is given in the steps of the compression algorithm below. The whole data set is represented as $V = \{\{I_1, P_{11}, P_{12}, \dots, P_{1p1}\}, \dots, \{I_N, P_{N1}, P_{N2}, \dots, P_{NpN}\}\}$ where I_i is an intra-coded frame in the i -th temporal group and P_{ij} is the j -th predictive frame in the i -th temporal group.

Assuming that there are only small changes between consecutive frames, wavelet transformation of changes instead of the entire frames yields higher compression ratios and lower decompression times. Therefore, the compression of an intraframe is independent of other frames while compression of a predictive frame is dependent on previous frames in the same temporal group. The overall compression algorithm is shown in figure 3. Note that compression is performed on each volume and only significant values contributing to the features are encoded. All 3D frames are decomposed into $4 \times 4 \times 4$ blocks and wavelet transformation is performed on each block contributing to the specified features in the volume. The steps of the compression algorithm are as follows:

1. **difference volume:** $\Delta V_k = V_k - V'_{k-1}$, where V_k is original image of k -th frame and V'_{k-1} is the reconstructed image of the compressed volume V_{k-1} . If V_k is an intra-coded frame, we assume $V'_{k-1} = 0$.
2. **wavelet transformation:** $W \Delta V_k$ = wavelet transformation of ΔV_k . Compute coefficients $c_1, \dots, c_m$ representing ΔV_k in a three-dimensional Haar wavelet basis.
3. **classification** Each wavelet coefficient c and block b is classified as either insignificant or significant. c and b are further classified based on which features they

contribute to. c and/or b can belong to more than one feature. In such cases, the survival of c or b is dependent on the highest weighted feature that contains them.

4. **truncation of blocks** A block which does not contribute to the features or has very small changes over time is considered as an insignificant block. By truncating insignificant blocks, we can achieve higher compression ratios and can control the time for the volume reconstruction. To identify blocks contributing to the i -th feature and having insignificant changes, the sum of the square of coefficients is compared with a threshold value λ_i . If the sum is less than λ_i , the block is truncated. For encoding the truncated block, only one bit is assigned in the block significance map.
5. **truncation of wavelet coefficients:** The i -th feature to be compressed has its own weight represented as a threshold value τ_i . By setting the threshold value, the reconstructed quality of a specific feature can be controlled. If a wavelet coefficient c associated with the i -th feature is less than τ_i , the coefficient is truncated into zero.
6. **encoding:** The overall encoding scheme is shown in figure 4. Once wavelet coefficient truncation is performed based on each features weight, we take the surviving coefficients and encode them. The encoding is performed on each block, resulting in a sequence of encoded blocks. We classify 64 coefficients in a block as one level-0 coefficient, 7 level-1 coefficients and 8×7 level-2 coefficients to take advantage of the hierarchical structure of a block.

In the header of a frame, a bit stream representing each block's significance is stored to indicate whether the block corresponding to each bit is a zero-block or not. This avoids additional storage overhead for insignificant blocks. One bit is assigned to each block in a sequence. Then, for each significant blocks, we store an 8bit map representing whether the one level-0 and seven level-1 coefficients are zero or not. Next, we have another 8 bit map representing whether each eight $2 \times 2 \times 2$ subblock has non-zero wavelet coefficients followed by a significance map for representing non-zero level-2 coefficients. After storing the level-2 coefficient significance maps, the actual values of non-zero wavelet coefficients are stored in order. We used two bytes for storing each coefficient value. Lossless compression is further applied to improve compression.

4.2 Seed Cells Insertion

To allow browsing clients to quickly extract isosurfaces encoded in a volume, seed cells are attached to the compressed stream of each frame. A seed cell is guaranteed to intersect with a connected component of the isosurface. By performing surface propagation from the given seed cells, we can avoid visiting unnecessary cells and save extraction time. Since only one seed cell per isosurface component is necessary, the size of the seed cells is negligible and search structures such as octrees and interval trees are not required. Therefore, the isosurface extraction time is only dependent on the number of triangles regardless of the volume size.

4.3 Run-Time Decompression

Since we have a sequence of wavelet encoded volumes $W \Delta V_k$, we can get an approximated image V'_k by decoding and performing an inverse wavelet transformation. More specifically, $\Delta V'_k = \text{inverse trans formation}(W \Delta V_k)$ and $V'_k = V'_{k-1} + \Delta V_k$. For the intra-coded frame V_i , $V'_{i-1} = 0$ and we can get V'_i directly from $\Delta V'_i$ with no dependency on other frames. Once V_i is reconstructed, succeeding predictive frames can be decoded frame by frame until the next intra-coded frame is reached.

The decompression is based on block-wise decoding. In intra-coded frames, every block needs to be decompressed with complete inverse wavelet transformation. On the other hand, in

predictive frames, only significantly changing blocks are updated so that it can approximate the actual image as accurately as possible while minimizing decompression time. The specific decoding algorithm is as follows. Using the block significance map, we can identify every significant block and its corresponding encoded blocks. For each encoded block, perform the following steps: Read 8bit $b_1^1, \dots, b_8^1$ to decide whether one level-0 coefficient and seven level-1 coefficient are zero or not. Next, read 8bit $b_1^2, \dots, b_8^2$ to decide whether each eight subblocks has non-zero value or not. If b_k^2 , where $k = 1, \dots, 8$, is set as 1, read 8bit $c_1^k, \dots, c_8^k$ to determine which coefficients of the k -th subblock are non-zero. From the significance maps read above, we can read the actual non-zero coefficients in order. When all values of the coefficients are determined, inverse transformation is performed to get the actual data values and the corresponding block is updated.

Once significant function values are decoded, we perform isosurface extraction. For each significant isovalue, we have a seed set, and hence we can perform the above extraction quickly.

5 Interactive Browsing

While compression ratio is an important factor for improving I/O performance in the memory and network systems, it is equally important that the visualization browser can read and interactively display compressed streams of multiresolution, time-varying data sets.

5.1 Multi-Resolution Isosurface Rendering

Since an isosurface often contains a lot of triangles, multiresolution techniques are necessary for saving both extraction and rendering time as a tradeoff with visual fidelity. One strength of wavelet transforms is that it provides multiresolution and compressed representations in a consistent format.

In our block-based wavelet transform, there are 3 levels consisting of one 0-th level, seven 1-st level, and fifty six 2-nd level coefficients. The 0-th level coefficient provides low-pass filtered average value of $4 \times 4 \times 4$ cells. 0-th level and 1-st level coefficients together provide approximated intermediate values averaging $2 \times 2 \times 2$ cells. When fast extraction and rendering of isosurfaces are more important than accuracy, the client can take only the 0-th level and/or 1-st level coefficients, and reconstruct a volume of lower resolution. Since the reconstructed low resolution volume is a good approximation of the original volume, not only is the extracted isosurface a good approximation of original isosurface, but the number of triangles extracted is also reduced. This process has the effect of low-pass filtering the volume, which can remove noise and artifacts incurred by lossy compression. Figure 6 shows three images rendered using the same volumetric data, but different levels of an isosurface.

5.2 Combined Rendering of Isosurface and Volumetric Data

We combine isosurface and volume rendering to take advantages of both techniques. We take an isosurface as the region we need to see in greater detail, with an opaque or translucent volume in the object space. The rendering results are shown and compared in figure 7. In this data set, isosurface extraction provides a shaded surface representing a specific temperature value. Volume rendering shows temperatures around the isosurface. By combining both techniques, visual information contained in the rendering is enhanced.

We tested our work on an implementation based on OpenGL. OpenGL gives us the ability to perform depth tests and maintain a depth map. We take advantage of this to render the isosurface and volume in a consistent manner. So we set the isosurface to be completely opaque and render it using OpenGL, with the depth test on. Then, we render a set of 3D textured polygons sliced

from a volume from back to front ordering. This is consistent with what is recommended in the OpenGL documentation [17].

During the rendering of the isosurface, we build up a depth map, which is used during volume rendering. While isocontouring in general needs a large amount of time, we already have the seeds required to perform the seed set isocontouring. Hence, we achieve fast isosurface reconstruction. When we perform volume rendering, to obtain correct transparency results, we render the polygons in a sorted order. While it is generally time consuming to perform polygon sort, Nvidia 3 graphics card's 3d texture mapping capability helps overcome this. The volume data is stored in the graphics card texture memory. We create polygons through a simple incremental algorithm, making slices through the volume. These polygons are rendered with the corresponding texture values in the 3D memory. An interactive transfer function map to control color and opacity values is used to obtain the required images. If the user rotates a volume, we need to update the slicing direction. To get slightly better performance, we turn off building the depth buffer (using `glDepthMask(0)`) when we render the volume, as we are sure of rendering the polygons in order, and since all polygons previously rendered are opaque.

6 Experimental Results

Compression and rendering results were computed on a PC equipped with a Pentium III 800MHz processor, 512MB main memory, and a NVidia GeForce 3 graphics card which has 128MB texture memory. We used standard OpenGL functions for 3D texture-mapping based volume rendering.

Table 1 provides information about our test data sets. The first data set is generated from a computational cosmology simulation for the formation of large-scale structures in the universe. Since the functions in the data set have negligible changes in the last few frames, we have given all our compression results, for this data set, based on the first 100 frames. The second data set is generated by an oceanographic simulation and represents temperature changes in the ocean over time. The original model has an approximate resolution of 1/6 degree (2160 by 960) in latitude and longitude and carries information at 30 depth levels. Since the original resolution of the data is too high for hardware volume rendering, we decimated it into $512 \times 256 \times 32$ by subsampling and took every third frame. The third data set is obtained from hemoglobin dynamics simulation. The simulation generated a sequence of hemoglobin pdb files with thirty time steps and electron density volumes are computed from each pdb file.

For testing the performance of our compression scheme, we encoded only high temperature regions ranging between 21.47 and 36.0 (celsius degree) in the oceanographic temperature data set and high density regions ranging between 0.23917 and 3.26161 in gas dynamics data set as shown in Figure 8 and 10. After encoding wavelet coefficients, we used *gzip* for lossless compression.

Table 2 and 3 show the reconstruction time, root mean squared error (RMSE), and the compression ratio changes over time in gas dynamics and oceanography data sets. The reconstruction time includes the time for disk read, *gunzip*, and decoding of wavelet coefficients. The RMSE was calculated using only those function values which contributed to the features. The compression ratio was calculated by comparing the size of original time-varying volume data and feature-based compressed data encoded by applying our lossy compression and *gzip*. As you can see in the tables, the reconstruction time of a P frame is much less than that of an I frame while the compression ratio of P frame is much higher than that of an I frame. The reason for this is that our compression scheme only encodes significantly changing blocks in P frames.

In figure 5, we compare our feature based encoding (FBE) scheme with the full volume encoding (FVE) scheme. Since FBE encodes only the values contributing to the specified features, both the decompression time and the compression ratio are significantly improved with respect to schemes encoding full volumes. While transmission and reconstruction times of volumetric and isosurface features are reduced by FBE encoding, client-side rendering is significantly accelerated by using PC graphics hardware.

Figures 8 and 9 show a typical frame of the gas dynamics data set compressed with different compression ratios. Figures 10 and 11 show the flow of two isosurfaces of specific temperatures. Figure 9 shows a zoomed view of the same set of volumes rendered in figure 8. We notice that good visual quality is maintained even at such zoom factors and high compression ratios. While a zoomed image of a volume compressed at a ratio of 148:1 is visually almost the same as the original, we get only a few artifacts at compression ratios of 301:1. Figures 10 and 11 show that similar results were obtained when our compression and rendering scheme was applied to the oceanographic data set. We have used isosurfaces to track the movement of water with a specific temperature and a translucent volumetric region to represent the surrounding temperatures. These figures demonstrate the strength of our scheme in being able to interactively render specific regions of interest with high quality isosurfaces, surrounded by related volumetric data. Timing results of rendering are presented in table 4.

Figure 12 shows the compression result of hemoglobin dynamics simulation. Volume rendering was applied to each compressed volumetric frame and we measured the frame rate. We achieved high compression ratio (110:1) and high frame rate (6.3 frame/sec) with reasonable reconstruction quality.

Generally, the frame rate of volumetric video playback is mostly dependent on the resolution of volumes and the number of triangles in the extracted isosurfaces. The size of rendering is also an important factor because the texture based volume rendering relies on per pixel operations.

Although wavelet based encoding can generate some losses in volumes as well as the topology in the reconstructed isosurfaces, we can achieve very high compression ratios with acceptable degradation.

7 Conclusion

We describe a lossy compression scheme for encoding time-varying isosurfaces with amorphous volumetric features specified by scalar value ranges. Since large time-varying volume data has significant coherence, compression is necessary for saving storage space, reducing transmission time, and improving the performance of visualizing the time-varying data. We have achieved several of our goals: (i) high compression ratio with minimal image degradation, (ii) fast decompression, by truncating insignificant blocks and wavelet coefficients, and (iii) interactive client-side rendering of compressed multiresolution isosurfaces and volumetric features. Therefore, our compression scheme is useful for the interactive navigation and exploration of time-varying isosurfaces with amorphous volumetric features residing in local and/or remote data servers.

Acknowledgments

An early version of this paper appeared in VolVis02 [25]. We are grateful to Anthony Thane developing the CCV volume rendering tool (Volume Rover) and library, and Xiaoyu Zhang and John Wiggins for writing the volumetric molecular blurring code. The cosmological simulations were performed by Hugo Martel, Paul R. Shapiro, and Marcelo Alvarez, of the Galaxy Formation and Intergalactic Medium Research Group at UT-Austin. The oceanographics simulation data was provided by Professor Detlef Stammer, Physical Oceanography Research Division of the Scripps

Institution of Oceanography, under the National Partnership for Advanced Computing Infrastructures project. The hemoglobin dynamics (pdb files) were provided by David Goodsell of Scripps Research Institute.

This work was supported in part by NSF grants ACI-9982297, CCR-9988357, a DOE-ASCI grant BD4485-MOID from LLNL/SNL and from grant UCSD 1018140 as part of NSF-NPACI, Interactive Environments Thrust, and a grant from Compaq for the 128 node PC cluster.

References

1. Bajaj, C.; Ihm, I.; Park, S. Visualization-specific compression of large volume data. Proc. of Pacific Graphics; Tokyo, Japan. 2001. p. 212-222.
2. Bajaj, Chandrajit; Ihm, Insung; Park, Sanghun. 3D RGB image compression for interactive applications. ACM Transactions on Graphics 2001;20(1):10-38.
3. Bajaj, Chandrajit; Pascucci, Valerio; Schikore, Daniel R. Fast isocontouring for improved interactivity. Proceedings of the 1996 Symposium for Volume Visualization 1996:39-46.
4. Cignoni P, Marino P, Montani E, Puppo E, Scopigno R. Speeding up isosurface extraction using interval trees. IEEE Transactions on Visualization and Computer Graphics 1997:158-170.
5. Ellsworth, David; Chiang, Ling-Jen; Shen, Han-Wei. Accelerating time-varying hardware volume rendering using tsp trees and color-based error metrics. Proceedings Volume Visualization and Graphics Symposium 2000:119-128.
6. Le Gall, Didier. MPEG: A video compression standard for multimedia applications. Communications of the ACM 1991;34(4):46-58.
7. Gerstner, Thomas; Pajarola, Renato. Topology preserving and controlled topology simplifying multiresolution isosurface extraction. In: Ertl, T.; Hamann, B.; Varshney, A., editors. Proceedings Visualization 2000. 2000. p. 259-266.
8. Guthe, Stefan; Straser, Wolfgang. Real-time decompression and visualization of animated volume data. IEEE Visualization 2001:349-356.
9. Guthe, Stefan; Wand, Michael; Gonser, Julius; Straßer, Wolfgang. Interactive rendering of large volume data sets. Proceedings of IEEE Visualization Conference 2002:54-60.
10. Howie CT, Black EH. The mesh propagation algorithm for isosurface construction. Computer Graphics Forum 1994;13(3):65-74.
11. Ihm I, Park S. Wavelet-based 3d compression scheme for interactive visualization of very large volume data. Proceedings of Graphics Interface 1998:107-116.
12. Khodakovsky, Andrei; Schröder, Peter; Sweldens, Wim. Progressive geometry compression. In: Akeley, Kurt, editor. Siggraph 2000, Computer Graphics Proceedings. ACM Press/ACM SIGGRAPH/Addison Wesley Longman; 2000. p. 271-278.
13. Lorensen WE, Cline HE. Marching cubes: A high resolution 3d surface construction algorithm. ACM SIGGRAPH 1987:163-169.
14. Lum, Eric B.; Ma, Kwan-Liu; Clyne, John. Texture hardware assisted rendering of time-varying volume data. IEEE Visualization 2001:263-270.
15. Ma, Kwan-Liu; Camp, David M. High performance visualization of time-varying volume data over a wide-area network status. Supercomputing. 2000
16. Machiraju, Raghu; Zhu, Zhifan; Fry, Bryan; Moorhead, Robert. Structure-significant representation of structured datasets. IEEE Transactions on Visualization and Computer Graphics 1998;4(2):117-132.
17. OpenGL. Developer faq. <http://www.opengl.org/developers/faqs/technical.html>
18. Pfister, Hanspeter; Lorensen, Bill; Bajaj, Chandrajit; Kindlmann, Gordon; Schroeder, Will; Avila, Lisa Sobierajski; Martin, Ken; Machiraju, Raghu; Lee, Jinho. The transfer function bake-off. IEEE Transactions on Computer Graphics and Applications 2001;21(3):16-22.
19. Said, Amir; Pearlman, William A. A new fast and efficient image codec based on set partitioning in hierarchical trees. IEEE Transactions on Circuits and Systems for Video Technology 1996;6:243-250.
20. Shamir, A.; Pascucci, V.; Bajaj, C. Proc of IEEE Visualization Conference 2000. Salt Lake City: Utah; 2000. Multi-resolution dynamic meshes with arbitrary deformations; p. 423-430.

21. Shapiro J. Embedded image coding using zerotrees of wavelet coefficients. *IEEE Transactions on Signal Processing* 1993;41(12):3445–3462.
22. Shen, Han-Wei. Isosurface extraction in time-varying fields using a temporal hierarchical index tree. *IEEE Visualization* 1998:159–166.
23. Shen, Han-Wei; Hansen, Charles D.; Livnat, Yarden; Johnson, Christopher R. Isosurfacing in span space with utmost efficiency (ISSUE). In: Yagel, Roni; Nielson, Gregory M., editors. *IEEE Visualization '96*. 1996. p. 287-294.
24. Silver D, Wang X. Tracking and visualization turbulent 3d features. *IEEE Transactions on Visualization and Computer Graphics* 1997;3(2):129–141.
25. Sohn, Bong-Soo; Bajaj, Chandrajit; Siddavanahalli, Vinay. Feature based volumetric video compression for interactive playback. *IEEE/SIGGRAPH Symposium on Volume Visualization and Graphics* 2002:89–96.
26. Sutton, Philip M.; Hansen, Charles D. Isosurface extraction in time-varying fields using a temporal branch-on-need tree (T-BON). In: Ebert, David; Gross, Markus; Hamann, Bernd, editors. *IEEE Visualization '99*. San Francisco: 1999. p. 147-154.
27. Taubin, Gabriel; Rossignac, Jarek. Geometric compression through topological surgery. *ACM Transactions on Graphics* 1998;17(2):84–115.
28. Westermann R, Thomas E. Efficiently using graphics hardware in volume rendering applications. *SIGGRAPH* 1998:169–177.
29. Zhang X, Bajaj C, Blanke W. Scalable isosurface visualization of massive datasets on cots clusters. *Proceedings of the IEEE Symposium on Parallel and Large-Data Visualization and Graphics* 2001:51–58.
30. Zhang, Xiaoyu; Bajaj, Chandrajit; Ramachandran, Vijaya. Parallel and out-of-core view-dependent isocontour visualization using random data distribution. *Joint Eurographics-IEEE TCVG Symposium on Visualization* 2002:9–18.

Fig. 1.
Interactive Volumetric Video Playback of Simulated Explosion Gas Dynamics.

Fig. 2.
Interactive Volumetric Video and Remote Streaming and Display Pipeline

Fig. 3.
Overall compression algorithm

Fig. 4.
Suggested encoding scheme for supporting fast decompression and high compression ratios

Fig. 5.

Performance comparison of different encoding schemes. Although the compression ratio of (a) and (b) is same (181:1), the quality of the reconstructed image using (a) is better than (b). RMSE of (a) is 0.110 and (b) is 0.131. (a) Only the values contributing to features are encoded. (b) Every value in the volume is encoded.

Fig. 6. Three different levels of an isosurface. The number of isosurface triangles and extraction time for level 0 (left), level 1 (middle), and level 2 (right) are (11878, 204ms), (41624, 625ms), and (207894, 3110ms).

Fig. 7.
The visualization of the oceanographic temperature change data set. Upperleft: Isosurface Rendering, Lowerleft: Volume Rendering, Right: Combined Rendering of Isosurface and Volume.

Fig. 8. Gas dynamics data set. Original volume (left), reconstructed volume with compression ratio 148:1 (middle), and compression ratio 301:1 (right).

Fig. 9.
Gas dynamics data set. Zoomed images from figure 8.

Fig. 10. Oceanographic temperature change data set. Original volume (left) and reconstructed volume with compression ratio 183:1 (right).

Fig. 11.
Oceanographic temperature change data set. Zoomed images from figure 10.

Fig. 12. Hemoglobin Dynamics Simulation. Left: Original Volume, Middle: Zoomed Original Volume, Right Zoomed Reconstructed Volume with Compression Ratio 110:1).

Table 1

Information on time-varying data sets

Data	Res.	type	#frm	1 frm size
Gas	$256 \times 256 \times 256$	float	144	64MB
Ocean(decim)	$512 \times 256 \times 32$	float	39	16MB
Hemoglobin	$128 \times 128 \times 128$	float	30	8MB

Table 2

Compression performance on gas dynamics data set (*: original density range [0., 8065.299])

Average Compression Ratio - 148:1, RMSE - 0.108			
Frame#	Recon. time	RMSE*	Comp. ratio
1(I)	687ms	0.105	40:1
2(P)	177ms	0.131	378:1
3(P)	271ms	0.101	182:1
4(P)	235ms	0.103	234:1

Average Compression Ratio - 301:1, RMSE - 0.139			
Frame#	Recon. time	RMSE*	Comp. ratio
1(I)	489ms	0.127	70:1
2(P)	141ms	0.156	988:1
3(P)	207ms	0.130	422:1
4(P)	188ms	0.134	516:1

Table 3

Compression performance on oceanographic data set (**: original temperature range [-2.0, 36.0])

Average Compression Ratio - 183:1, RMSE - 0.090			
Frame#	Recon. time	RMSE**	Comp. ratio
1(I)	124ms	0.076	72:1
2(P)	76ms	0.087	273:1
3(P)	96ms	0.089	226:1
4(P)	86ms	0.091	223:1

Average Compression Ratio - 348:1, RMSE - 0.177			
Frame#	Recon. time	RMSE**	Comp. Ratio
1(I)	106ms	0.157	103:1
2(P)	57ms	0.169	615:1
3(P)	65ms	0.175	493:1
4(P)	64ms	0.178	482:1

Table 4

Timing results of rendering isosurfaces with amorphous volumetric features in one frame:(the data set name, 3D texture loading time, isosurface extraction time, number of triangles in the isosurfaces, isosurface rendering time, volume rendering time)

Data set	Load	Ext. Time	Tri#	Isosurface	Volume
Gas	701ms	1703ms	135362	312ms	422ms
Ocean	156ms	1640ms	104900	235ms	281ms