

Memory Efficient Acceleration Structures and Techniques for CPU-based Volume Raycasting of Large Data

Sören Grimm, Stefan Bruckner, Armin Kanitsar and Eduard Gröller[†]

Institute of Computer Graphics and Algorithms
Vienna University of Technology

Abstract

Most CPU-based volume raycasting approaches achieve high performance by advanced memory layouts, space subdivision, and excessive pre-computing. Such approaches typically need an enormous amount of memory. They are limited to sizes which do not satisfy the medical data used in daily clinical routine. We present a new volume raycasting approach based on image-ordered raycasting with object-ordered processing, which is able to perform high-quality rendering of very large medical data in real-time on commodity computers. For large medical data such as the Visible Male (587x341x1878) we achieve rendering times up to 2.5 fps on a commodity notebook. We achieve this by introducing a memory efficient acceleration technique for on-the-fly gradient estimation and a memory efficient hybrid removal and skipping technique of transparent regions. We employ quantized binary histograms, granular resolution octrees, and a cell invisibility cache. These acceleration structures require a small extra storage of approximately 10%.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Three-Dimensional graphics and realism, Raytracing, Visible line/surface algorithms.

1. Introduction

Direct Volume Rendering is known as a powerful technique to visualize complex structures within three-dimensional data. Its main advantage, compared to standard 3D surface rendering, is the ability to perform translucent rendering in order to provide more information about spatial relationships of different structures. In general 3D visualization helps to understand patient's pathological conditions, improves surgical planning, and is a big aid in medical education. However, a typical data size of today's clinical routine is about 512x512x1024 (16 bit CT data) and will increase dramatically in the near future. This presents a challenge to current rendering architectures and techniques. The increasing demand of interactive 3D visualization is basically driven by the size itself. Conventional slicing methods already reach their limit of usability due to the enormous amount of slices. And so 3D visualization becomes more and more an attrac-

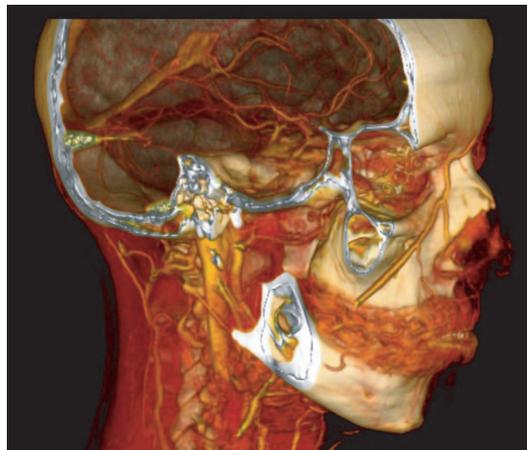


Figure 1: CTA scan of human head with enhanced venous system.

[†] {grimm, bruckner, kanitsar, groeller}@cg.tuwien.ac.at. Favoritenstrasse 9-11, A-1040 Vienna, Austria.

tive alternative method for examinations of large medical data. Figure 1 shows an example of 3D visualization.

Within the research area of accelerating volume rendering, two main research streams may be distinguished. One stream is focused on exploiting special purpose hardware such as Volume Pro [PHK*99], Vizard [MKW*02] or graphic cards (GPU) [WE98, GWGS02]. This approach usually provides high performance when data fits into internal memory. However, this issue becomes the most critical bottleneck once the data size exceeds the onboard internal memory capacity. Expensive main memory to internal memory transfers have to be performed, which lead to an enormous performance penalty. Furthermore, the accelerated pace of the GPUs development cycle forces the heterogeneity of multi-user hardware environments. This makes the adoption of such special purpose hardware solutions even more difficult.

The other research stream is based on CPU technologies. In general they provide better performance for large data due to the inherent larger memory capacity. Many proposed approaches for CPU based volume raycasting achieve high performance by utilizing super-computers or clusters; e.g. Parker et al. [PSL*98] presented a volume rendering approach on an SGI Reality Monster and was capable to render the Visible Woman (approx. 1 GB) with up to 20 fps utilizing 128 processors. However, it is a large scale solution which does not apply to the needs and capacities of an ordinary medical environment.

The purpose of this paper is to present a solution which resolves the issues presented before: an interactive real-time volume rendering approach for large medical data, capable of performing in a heterogeneous hardware environment, by using commodity computers such as notebooks, and providing high performance and high quality images. We achieve this by introducing an efficient method for on-the-fly gradient estimation and an efficient hybrid removal and skipping technique of transparent regions. The presentation of our new approaches is subdivided as follows: Section 2 surveys related work. Section 3 presents a brief overview of our raycasting processing work-flow. In section 4 we introduce our refined caching scheme to accelerate on-the-fly gradient estimation. In section 5 we focus on removing and skipping of transparent regions by employing quantized binary histograms, granular resolution octrees, and a cell invisibility cache. In section 6 we discuss and conclude our work. Finally in section 7 we give ideas for future work.

2. Related Work

The most popular CPU-based direct volume rendering algorithms are shear-warp, splatting, and raycasting. Shear-warp is considered to be the fastest software algorithm [LL94], however the inherent bi-linear interpolation provides quality which is in general insufficient for medical purposes. Splatting was first proposed by Westover et al. [Wes90]. Later

it was improved in terms of quality and speed by Mueller et al. [MMC99, MSHC99] and Huang et al. [HMSC00]. This technique provides high quality images. However it still lacks the speed provided by the general volume raycasting technique.

Volume raycasting is still widely used if high quality rendering of large data is desired. Several acceleration techniques for volume raycasting have been proposed over the last decade. Knittel [Kni00] and Mora et al. [MJC02] proposed volume raycasting approaches for commodity computers. They achieve impressive frame-rates by using a spread memory layout and pre-computed gradients; however their method requires a huge amount of additional memory. The spread memory layout itself increases the memory usage by a factor of four. This becomes a rather limitation factor if large data needs to be handled, or if the rendering system is part of a larger visualization systems and memory resources need to be shared.

In contrast to that, our approach does not rely on pre-computing or a spread memory layout; it is based on a bricked volume layout. However, to achieve high performance advanced acceleration structures and techniques are necessary. In the following sections we present several memory efficient acceleration approaches.

3. Volume Raycasting Work-flow

The following paragraph presents a brief overview of the work-flow of our volume raycasting approach. Bricking of volume data is a well known method to exploit cache coherence [PSL*98, GWGS02, LY96, GBKG]. We decompose the volume data in bricks and perform processing brick-wise. The volume raycasting process is subdivided into pre-processing, pre-rendering, rendering, and post-rendering. The pre-processing step is done only once during start-up and the remaining steps are performed every time the image needs to be re-rendered. At first we give a brief overview of the four rendering steps.

Pre-Processing: During loading, the data is decomposed into small bricks of size 32^3 . The data within the bricks and the bricks themselves are stored in common xyz-order. For each brick information about the contained density values is stored, e.g. min-max values, quantized binary histograms, etc.

Pre-rendering: In this phase transparent regions are removed and the rays-volume intersections are computed. There are eight different brick processing lists which are defined by the eight possible viewing-quadrants in 3D. Depending on the viewing direction the appropriate list is selected to process the volume brick-wise and in correct visibility order.

Rendering: According to the brick list, all rays traverse the bricks in visibility order, until all bricks are processed or all rays are terminated due to complete opacity accumula-

tion. During traversing regular re-sampling, gradient computation, classification, shading and composition are performed.

Post-rendering: At this point the final image is displayed with OpenGL or similar painting methods, or written to a file, or sent over the network to a client.

A more detailed description of the used acceleration techniques and structures for the pre-rendering and rendering step is given in the following sections. There are two major strategies to accelerate volume raycasting. The first one is to reduce the computational costs at one re-sampling location. We achieve this by using an acceleration technique for gradient estimation (section 4). The second is to efficiently remove and skip transparent regions, which we achieve by using quantized binary histograms, granular resolution occlusion trees, and a cell invisibility cache (section 5).

4. Efficient Gradient Caching

The most common method to accelerate gradient estimation is to read pre-computed gradients from memory. However, this acceleration technique has several drawbacks. In order to gain high performance the gradients must be stored in memory, resulting into an inefficient usage of resources. Furthermore such a solution is limited by memory bandwidth instead of the preferable CPU throughput; the evolution of computer systems has shown that CPU performance increases faster than memory bandwidth. Going one step further if the data exceeds the main memory capacity, out-of-core rendering has to be performed and the gap between CPU throughput and memory bandwidth becomes even larger. Additionally, experience has shown, that not every gradient estimation scheme performs equally well on all kinds of data. Therefore the ability to switch between different gradient estimation schemes is an important feature and basically not efficiently given if pre-computing is used. Finally, pre-computing the gradients is quite time consuming. Considering a now-a-days medical visualization system, the doctors main interest is to carry out the examination as fast as possible. The total time from scanning the patient to the actual examination is a highly critical factor.

To avoid these issues, our approach performs on-the-fly gradient estimation. However, in order to obtain highly accurate images, a dense object and image sample distance is inevitable, which implies high computational costs. A typical re-sampling resolution illustrated in 2D is shown in figure 2a. In this case there are eight re-sample locations within a cell. Each gradient at the corners of one cell has to be computed eight times. Furthermore, each corner is shared between four cells in 2D. The total amount of redundant gradient computations at one corner add up to eight re-sampling positions multiplied by four cells which gives a total of 32 computations. Mapping the same calculation to 3D, results in computational costs which are considerably higher. These very costly redundant gradient computations can be avoided

by refined caching. However, not every gradient estimation scheme is suitable for caching. There are several studies on gradient filters for volume rendering with focus on accuracy, importance in terms of image quality and efficiency. Especially, Möller et al. [MMMY97] give a thorough comparison of commonly used normal estimation schemes. They differentiate between four types of gradient estimation schemes:

1. *Continuous Derivative* uses a derivative filter which is pre-convolved with the interpolation filter. The gradient at a resample location is then computed by convolving the volume by this combined filter.
2. *Analytic Derivative* uses a special gradient filter derived from the interpolation filter for gradient estimation.
3. *Interpolation First* computes the derivative at the resample position by re-sampling the data on a new grid, such that the used derivative operator can directly be applied. This is very beneficial if orthographic rendering is performed.
4. *Derivative First* computes the gradients at the grid-points and then interpolates these at the desired re-sample position.

For scheme one and two no caching mechanism is available. Only schemes three and four can be considered for gradient estimation. Due to their numerical equivalency only a comparison with respect to efficiency is necessary. Möller et al. [MMMY97] proposed the *Interpolation First* method as the most efficient one. Considering volume rendering and no caching this is quite obvious. However, applying the *Interpolation First* scheme requires re-sampling of the original grid to a much larger grid if the object sample distance is significantly smaller than one. Already an object sample distance of 0.25 increases the grid size by a factor of four. This enormous amount of data makes caching inefficient and difficult. Especially if the object sample distance should be kept dynamical or if jittering techniques to improve the image accuracy are applied. Furthermore if perspective rendering is required, caching becomes impossible. Due to these reasons the *Derivative First* gradient estimation scheme is more efficient from a performance point of view, since it is more suitable for caching. In this case, the amount of data to cache is always determined. This makes interactive changes of the object sample distance possible. In addition perspective rendering is also accelerated with the same caching mechanism.

4.1. Per Brick Gradient Caching

Our caching scheme requires two data structures: the cache itself and a second structure to store the corresponding valid bits. The used processing entity is not the whole volume; in fact the volume is decomposed in bricks and each brick defines a processing entity. The size of the cache matches the number of gradients needed for one brick. The most straightforward way to use this cache would be to pre-compute all gradients which correspond to the current brick and use

those during brick processing. This would be very inefficient, since more gradients than necessary would be pre-computed if only parts of a brick are visible. In contrast to that we additionally use valid bits, which encode if a gradient is already computed and stored in the cache. During brick processing every time a gradient needs to be computed, it is checked if the gradient is already stored in the gradient cache. If not, the gradient is computed and stored in the cache and the corresponding valid bit is set to true. This mechanism ensures that gradients are computed only once at each sample position during brick processing. The cache remains only valid during the processing of one brick. Once the next brick is processed the cache is reset. This has the effect that the gradients which are also needed in adjacent bricks are processed more than once. The resulting performance penalty is low, since the number of those gradients is small compared to the number of all gradients.

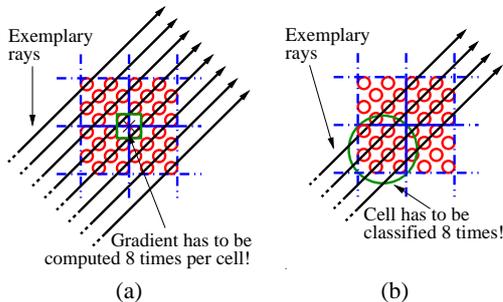


Figure 2: Typical re-sampling resolution of a cell in 2D. (a) In the shown case each gradient at the cell corners has to be computed 8 times while processing one cell. (b) In the shown case a cell has to be classified 8 times.

4.2. Performance Results and Memory Consumption

The memory consumption of the cache is rather low. The cache size is not related to the volume dimensions. It is related to the brick dimensions. The brick dimensions in our case are $32 \times 32 \times 32$, the size of the gradient cache is $(dimension\ of\ brick + 1)^3$ multiplied by $dimension\ of\ gradient$ multiplied by $size\ of\ gradient\ component$, which is $(33)^3 \cdot 3 \cdot 4 \approx 421,14$ KByte. Additionally we store for each cache entry a valid bit, which adds up to $33^3 \approx 4.39$ KByte. This is altogether less than 512 KB. For performance reasons the data shall remain in the level 2 cache. This is not an issue as current commodity CPUs have level 2 cache size of 1 MB.

Figure 3 shows the effect of per brick gradient caching compared to per cell gradient caching and no gradient caching at all. Per cell gradient caching means that gradients are cached while a ray re-samples a cell. For gradient estimation we used the gradient filter proposed by Neumann et al. [NCV*02]. This filter produces slightly better quality than the Sobel filter, supports inherent volume filtering

and has approximately the same computational costs. Due to the on-the-fly computations, the filtering can be enabled and disabled interactively. The on-the-fly filtering has low computational costs and can be used to increase the quality, when less number of rays are shot to increase the frame-rate during interaction.

Our testing system was an Intel®Pentium®M 1.6 GHz with 1 MB level 2 cache and the used data was the UNC head. We chose an adequate opacity transfer function to enforce translucent rendering. The charts from left to right show different timings for object sample distances from 1.0 to 0.125 for three different zooming factors 0.5, 1.0, and 2.0. In case of zooming factor 1.0 we have one ray per cell, already here per brick gradient caching performs better than per cell gradient caching. This is due to the shared gradients between cells. For zooming out (0.5) both gradient caching schemes perform equally well. The rays are as far apart such that nearly no gradients can be shared. On the other hand for zooming in (2.0), per brick caching performs much better than per cell caching. This is due to the increased number of rays per cell. As more rays process the same cell, as more beneficial the per brick caching becomes. Per brick gradient caching compared to no caching shows already with a zoom factor of 2.0 and an object sample distance of 0.5 an impressive speedup of approximately 3.0. The speedup favorably scales with the zoom factor.

5. Removing and Skipping of Transparent Regions

For medical imaging, interactive classification of data is mandatory. In general during examination it happens quite often, that large parts of the data are classified as transparent to allow a more precise view of the region of interest. For acceleration purposes it is quite beneficial to exploit this transparency information and start the actual re-sampling of the data right where the visible data begins. The work-flow of our hybrid transparent region removal and skipping technique is shown in figure 4. At first transparent regions are removed on a brick basis (figure 4a -> figure 4b). Then to support even more refined removal of smaller transparent regions we perform octree projection (figure 4b -> figure 4c). Due to efficiency reasons our octree subdivision does not fully go down to individual cells. The granular resolution of the octree leads to approximate rays-volume intersections. To overcome the resulting performance penalty we introduce a Cell Invisibility Cache (CIC) to skip the remaining transparent cells (figure 4c -> figure 4d). In the following we describe our hybrid transparent region removal and skipping technique in more detail.

5.1. Quantized Binary Histograms

At first we describe an efficient encoding for finding transparent bricks. The most common method are minimum-maximum value encodings and a summed area table eval-

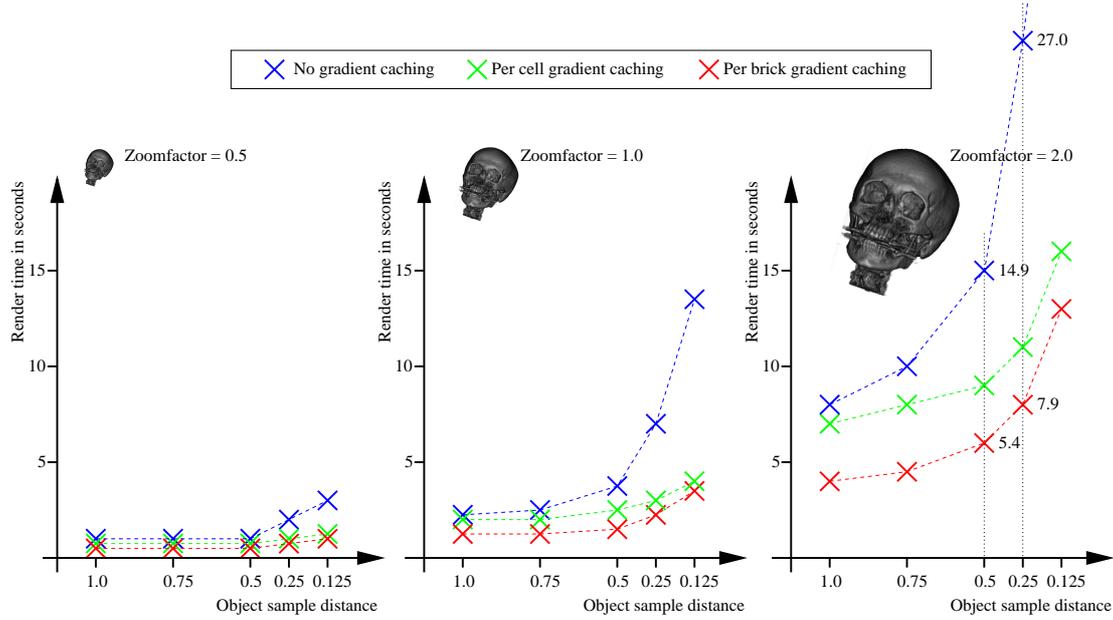


Figure 3: Comparison between no gradient caching, per cell gradient caching, and per brick gradient caching. The used data was the UNC head (256x256x224 - 12 bit), intensity range [0,1136] was mapped to 0.0 opacity and range [1136,4095] to a linear opacity ramp between 0.0 and 1.0. Hardware: CPU - Intel®Pentium®M 1.6 GHz, Cache - 1 MB Level2, RAM - 1 GB, GPU - GForce4 4200 Go (32MB).

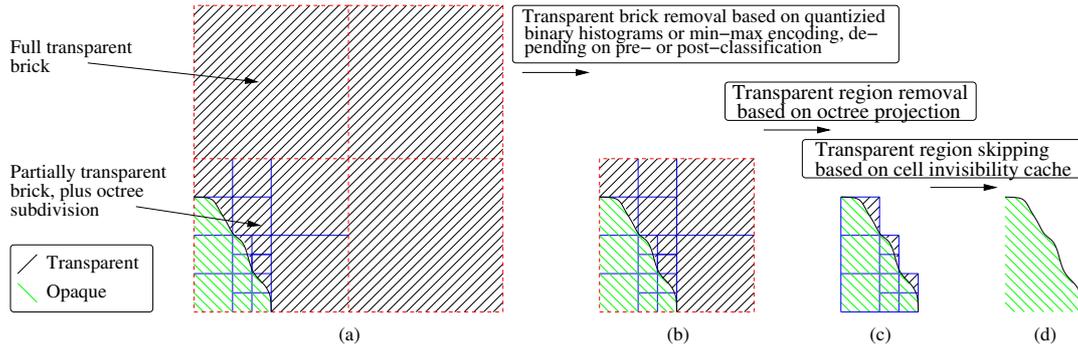


Figure 4: General work-flow of our hybrid transparent region removal and skipping technique.

uation. A summed area table encodes the opacity integral by

$$S(0) = \alpha(0)$$

$$S(k) = S(k-1) + \alpha(k)$$

Hereby $k \in H = [0..4095]$, which is the possible range of Hounsfield units and α represents the opacity. i_{min} and i_{max} denote the minimum and maximum density value within a brick. The integral of the discrete function α over the interval $[i_{min}, i_{max}]$ can be evaluated in constant time by performing two table lookups:

$$\sum_{k=i_{min}}^{i_{max}} \alpha(k) = S(i_{max}) - S(i_{min})$$

submitted to Symposium on Visualization (2004)

If $S(i_{max}) - S(i_{min}) = 0$ then the brick is transparent and can be skipped. At this point we differentiate between pre- and post-classification. For post-classification the min-max encoding is the most accurate, since due to interpolation of data all values between the minimum and the maximum may occur. However, if pre-classification is performed the min-max encoding may be too granular when applied on large regions. Figure 5 shows an example where the min-max encoding fails if the region of interest is set 150. The min-max encoding would report both bricks as being visible. The main issue is, that the min-max encoding accuracy relies heavily on the underlying data. If the region is large it is quite likely that its

values differ considerably. The min-max encoding becomes too granular to effectively encode the area. We use a more refined structure, i.e. a quantized binary histogram. In general a binary histogram is encoded as:

$$\sigma_x(B) = \begin{cases} 1, & x \in B \\ 0, & \text{otherwise} \end{cases}$$

Hereby B is the set of all density values a brick contains, with $B \subseteq H = [0..4095]$. $\sigma_x(B) = 1$ means that the density value x is given at least at one grid position in the underlying brick. This encoding is effective, however it is quite inefficient in terms of memory usage and efficient evaluation. We store instead of the full binary histogram a quantized binary histogram. It is determined by

$$\sigma_i(B) = \begin{cases} 1, & \forall x \in [128 \cdot i..128 \cdot (i+1)] [x \in B, (0 \leq i \leq 31)] \\ 0, & \text{otherwise} \end{cases}$$

Within quantized binary histograms the existence of data within a specific interval is encoded. The intervals are concatenated, disjunct, have same length, and cover the range of Houndsfield units. In the pre-processing phase every brick is parsed and encoded. The same encoding can be performed for the transfer-function with respect to opacity:

$$\lambda_i = \begin{cases} 1, & \forall x \in [128 \cdot i..128 \cdot (i+1)] [opacity(x) \neq 0] \\ 0, & \text{otherwise} \end{cases}$$

Hereby $x \in H$ and $i \in [0,31]$. Every time the transfer-function changes, the transfer-function is re-encoded in this way.

With this information one can quickly determine the transparent bricks. A brick is transparent if

$$\forall i \in [0..31] : \lambda_i \wedge \sigma_i = 0$$

This conjunction test can be done very efficiently on a x86 based CPU. Note, that this is a conservative estimate of bricks visibility. It is possible that due to the chosen encoding we consider a brick visible although all contained values are classified as transparent. However, if we look at figure 5 and set the region of interest to 150, we can see that the quantized binary histogram would report the bricks correctly if pre-classification is performed. This is due to the fact, that the quantized binary histogram is more sensitive for largely varying data values. This property can be efficiently exploited if the binary histogram encodes a segmentation information volume. In such a volume, segmented objects are encoded by labels. These labels can differ largely and interpolation is not applicable, so only pre-classification can be performed.

5.2. Granular Resolution Octrees

After eliminating entire transparent bricks with quantized binary histograms, we use granular resolution octrees to eliminate transparent regions of a single brick. If it comes to

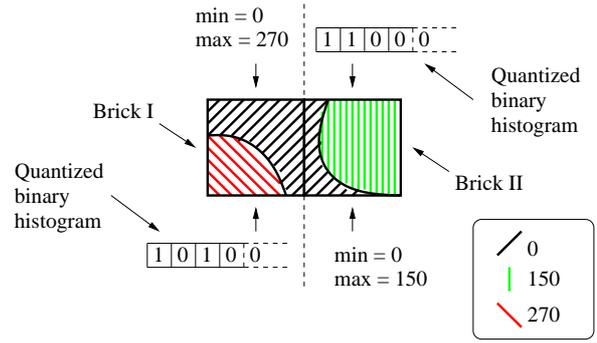


Figure 5: Min-Max encoding granularity issue if pre-classification is performed: If the chosen visible area is 150, the min-max encoding would report both bricks being visible. Quantized binary histograms would report Brick I being invisible and Brick II being visible.

small spatial regions one of the best known space subdivision structures to support refined skipping of transparent areas is a min-max octree [LL94, WG92, MJC02]. Each brick (32x32x32) contains a 3-level min-max octree, shown in figure 6a. For each octree level we store the minimum and maximum value as a pair of numbers. For level 0 we have 8 pairs, level 1 needs 8x8 = 64 pairs, and level 2 needs 8x8x8 = 512 pairs. When classification changes the octree is recursively evaluated by a summed area table for all bricks. We store the classification information efficiently by hierarchical compression. Nodes of level 2 are either opaque or transparent. All other nodes have an additional inhomogeneous state. The information whether a node of level 2 is transparent or opaque is stored in one bit. The state of a level 1 node is determined by testing of one byte, which contains all the bits of its children. For level 0 such a hierarchical compression requires to test 8 bytes for a node and 64 bytes for the brick. Due to efficiency reasons we additionally explicitly store the state information of level 0. We have three possible states, thus we need 2 bits for each level 0 node. For each node we store one byte for the transparency information and one byte for the homogenous information. A bit in the transparency byte is only defined if the corresponding bit in the homogenous byte states a node as homogenous. Due to this encoding, the octree can be very efficiently traversed.

5.3. Removing of Transparent Regions

We have two structures, a quantized binary histogram and a granular octree, to find the rays-volume intersections up to the resolution of the granular octree (figure 4c). The bricked geometry of the volume and the octrees within the bricks are converted to a polygonal structure and rendered into a z-buffer. Basically we traverse through the brick list and determine which brick has visible data and needs to be evaluated for rays-volume intersection. The evaluation is performed by

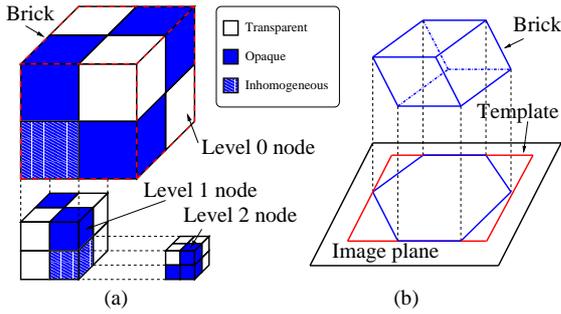


Figure 6: (a) Octree classification scheme for an individual brick. Sizes - Brick: $32 \times 32 \times 32$, Level 0 node: $16 \times 16 \times 16$, Level 1 node: $8 \times 8 \times 8$, and Level 2 node: $4 \times 4 \times 4$. (b) Brick projection template.

quantized binary histograms in case of pre-classification or min-max encodings in case of post-classification. The octree of those bricks is evaluated and the sub-bricks which contain visible data are rendered. This rays-volume intersection computation by rendering requires the *granular* resolution of the octree. Otherwise the number of polygons would exceed the rendering performance of commodity graphics hardware.

Utilizing OpenGL for rendering provides high performance and high accuracy; however, if the approach is used as an integrated module it requires off-screen rendering. This is available in OpenGL by PBuffers. Unfortunately, this feature is not available on every graphics card. Furthermore the rendering requires a huge amount of graphics cards memory. Considering a 1024×1024 image, the needed buffer is already 8 MB. Most of the more advanced medical visualization systems support high-resolution dual-displays. This feature normally utilizes all the available graphics card memory. There is no space left for graphics hardware accelerated off-screen rendering. Due to this reason, we also developed the rendering in software. This can be done very efficient, if the simple polygonal structure of the bricked octree layout is exploited. Since every brick is of the same structure, one can use template based projection of the octree. Similar work has been done by Srinivasan et al [SFH97]. The main idea is to project just one brick per viewing direction for each octree level as shown in figure 6b. This projection is used as a template for all other bricks of the same level. Any other brick of the same level has the same projection footprint and is obtained by translation. The projected footprint actually consists of z-values, since we are interested in the z-buffer footprint of the octree. All possible entry bricks are rendered in a front-to-back order by using the projected z-value template. The resulting z-buffer footprint of the octree is then used to determine the rays-volume intersections. This is as fast as the OpenGL implementation, since the costly projection itself has to be done only for one brick per viewing direction.

Furthermore no costly OpenGL `glReadBuffer()` instruction is involved and the resulting z-buffer contains directly the z-components of the ray starting-positions.

5.4. Cell Invisibility Cache: Skipping of Transparent Regions

As the granular resolution octree does not go down to cell level, a cell invisibility cache is used to skip the remaining transparent cells (figure 4c -> figure 4d). The volume-rays intersections estimation by template-based projection of the octree subbricks brings us as close as $4 \times 4 \times 4$ samples to the visible data. This is inefficient from a performance point of view. Especially if first-hit-raycasting is performed every non skipped sample has a large impact on the resulting frame-rate. A resolution of $4 \times 4 \times 4$ results in a large number of non skipped samples, this is depicted by the red samples shown in figure 8. All these samples have to be classified in order to determine which cell can be skipped. Depending on the object-sample distance and the zoom factor these cells have to be classified several times. This is shown for a typical re-sampling resolution in figure 2b. In this case each cell has to be classified eight times. Considering the same example in 3D, the number of redundant cell classifications would be considerable larger. Due to this reason we introduce refined cell invisibility caching. We extend the volume raycasting pipeline in such way that classification of these invisible cells has to be done only once. The extended pipeline is shown in figure 7. A Cell Invisibility Cache (CIC) is attached at the beginning of the traditional volume raycasting pipeline. This CIC is initialized in such a way that it reports every cell as visible. In other words every cell has to be classified. Now, if a ray is sent down the pipeline, every time a cell is classified invisible this information is cached in the CIC. A cell can either be invisible or visible, this information can be encoded in just one bit. Once a cell is classified as invisible, the costly classification of a whole cell is exchanged by a binary test. This leads to an enormous performance increase. On the one hand, due to the reduced memory access and on the other hand due to the inherent classification and conjunction information of 8 samples. The information stored in CIC remains valid as long no transfer-function change is performed. During the examination of the data, e.g. by changing the viewing direction, the CIC fills up and the performance increases progressively. The same mechanism is also very beneficial for general empty space skipping within the data.

5.5. Performance Results and Memory Consumption

The additional memory usage of all three acceleration structures, quantized binary histogram, granular resolution octree, and the cell invisibility cache is rather low. Considering the size of the volume as 100%, they increase the size by approximately 10%. We use bricks of size $32 \times 32 \times 32$ stor-

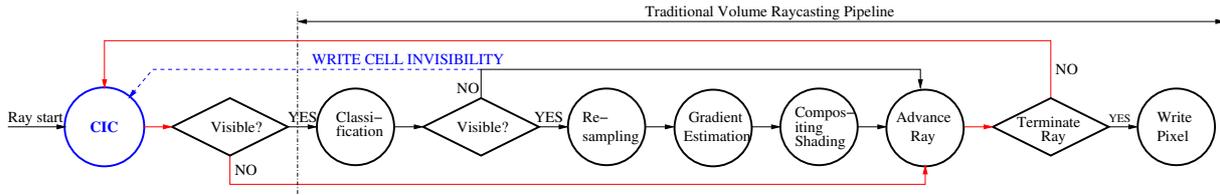


Figure 7: Cell Invisibility Cache (CIC) - Acceleration by caching invisibility information of cells. The acceleration path is emphasized in red.

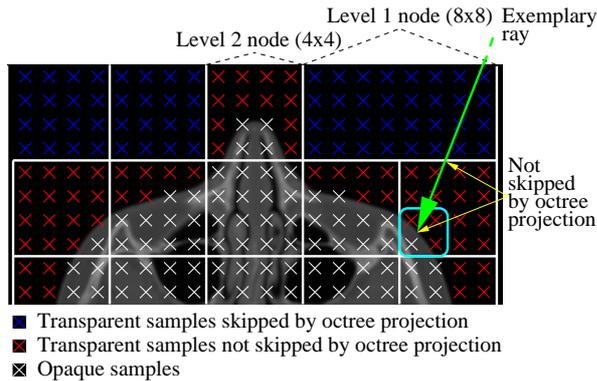


Figure 8: Zoomed in granular octree of one brick. The red samples can not be skipped due to the granular resolution of the octree.

ing 2 bytes for each sample, which is a total of 65536 bytes. Additionally for each brick we store:

- Quantized binary histogram: 4 byte
- Min-max information: $(512+64+8+1) * 4 = 2340$ byte
- Octree classification information: $(64 + 2) = 66$ byte
- Cell Invisibility Cache: $32^3 \text{ bit} / 8 = 4096$ byte

In total the storage increase is $((4 + 2340 + 66 + 4096) / 65536) \cdot 100 \approx 9.9\%$.

Table 1 shows the effect of our hybrid removal and skipping technique of transparent regions and figure 9 shows the corresponding rendering output. For benchmarking we used a commodity notebook equipped with an Intel®Pentium®M 1.6 GHz CPU, 1 MB Level2 cache, 1 GB RAM, and a GeForce4 4200 Go (32MB). The graphics card capabilities are only used to display the final image. We tested four different data sets. A rather small data set, the UNC head to be able to compare our speed to the approach of Mora et al. [MJC02]. This approach is slightly faster than the UltraVis system [Kni00]. They are both based on a spread memory layout and use pre-computed gradients. This leads to an inefficient memory usage and so they are restricted to rather small data. Mora's total render time is approximately a factor of two faster than our approach. However, Mora's approach uses pre-computed gradients, does pre-shading, and

its template based interpolation scheme limits the zooming to a zooming-factor of four. In contrast to that we chose to sacrifice some performance for increased flexibility, high quality, and a significantly lower memory usage. This enables us to render large data, used in daily clinical routine, on commodity hardware. We tested three different large typical medical data sets. The results show that our acceleration techniques typically achieve render-times of about 2 fps even for these large data sets. Table 1, fourth column, shows the total render time achieved by brick-wise transparent region removal. In the fifth column we additionally apply the granular octree projection. And finally in the sixth column we enabled the Cell Invisibility Cache to see the overall total render time achieved by the combined effect of all three acceleration structures.

6. Discussion and Conclusion

We presented a volume raycasting approach which provides high-quality images in real-time for large data on standard commodity computers without advanced graphics hardware. For large medical data such as the Visible Male (587x341x1878) we achieve rendering times up to 2.5 fps on a commodity notebook. Furthermore, our approach can utilize symmetric multiprocessing systems as processing is performed brick-wise. It scales well and achieves a speedup factor of approximately 2.0 on a dual CPU machine. This is very beneficial if a large amount of data has to be processed. Although we avoided any pre-computing and compute each part of the volume raycasting pipeline on-the-fly, we achieved performance in the same range as approaches which heavily rely on the memory bandwidth [MJC02]. Our refined caching scheme for gradient estimation in conjunction with hybrid skipping and removal of transparent regions, enables us to achieve high quality while maintaining high performance. Due to the efficient memory consumption of our acceleration structures (quantized binary histogram + granular resolution octree + Cell Invisibility Cache) and the bricked volume layout we are able to handle very large data. All acceleration structures require only an extra storage of approximately 10%. Data sizes up to 2GB are possible, a limitation imposed by the virtual address space of current consumer operating systems.

Name	Dimensions	Size	Transparent brick removal	Octree projection	CIC / Total render time	Mora
(a) Visible Male	587x341x1878	0.70 GB	0.61 sec	0.46 sec	0.40 sec	NA*
(b) Visible Male	587x341x1878	0.70 GB	0.68 sec	0.53 sec	0.45 sec	NA*
(f) Aorta	512x512x1112	0.54 GB	1.16 sec	0.93 sec	0.61 sec	NA*
(d) Lower extremities	512x512x1202	0.59 GB	0.86 sec	0.70 sec	0.64 sec	NA*
(e) Lower extremities	512x512x1202	0.59 GB	0.69 sec	0.46 sec	0.37 sec	NA*
(c) UNC head	256x256x256	0.03 GB	0.71 sec	0.26 sec	0.18 sec	0.075 sec

Table 1: Performance results for different data sizes, which are used in daily clinical routine. Image size: 512x512, Sample rate: 0.5, and Hardware: CPU - Intel®Pentium®M 1.6 GHz, Cache - 1 MB Level2, RAM - 1 GB, GPU - GForce4 4200 Go (32MB). * Unable to load data set, system memory exceeded.

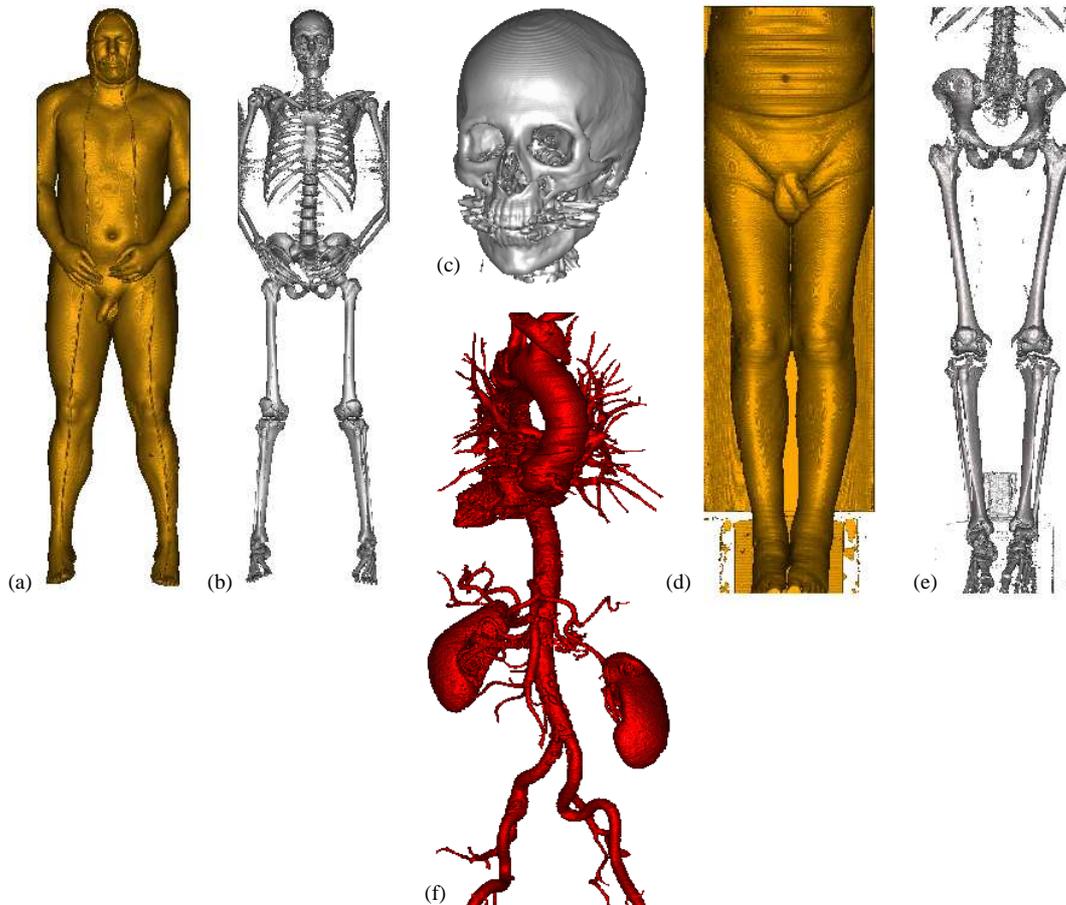


Figure 9: Result images of benchmarking.

7. Future Work

In the future we want to support out-of-core rendering to be prepared for the next generation of data sizes. First commodity prototype scanners already deliver data in the range of 1024x1024x2048, which is about 4GB of data and quite

challenging to handle. Additionally we want to incorporate brick-based compression to achieve interactive rendering times for these large data. Furthermore we want to support perspective rendering. The changes which have to be made, are basically to split the brick rendering lists such

that the correct processing order is ensured. Additionally we have to render perspective distorted bricks to be able to find the exact entry points of the rays. Other than that all the presented accelerating techniques should work as well as for parallel projection.

Additional material (images, movies):
<http://www.cg.tuwien.ac.at/~sgrimm/vissym2004/>

8. Acknowledgements

The work presented in this publication has been funded by the ADAPT project (FFF-804544). ADAPT is supported by *Tiani Medgraph*, Vienna (<http://www.tiani.com>), and the *Forschungsförderungsfonds für die gewerbliche Wirtschaft*, Austria. See <http://www.cg.tuwien.ac.at/research/vis/adapt> for further information on this project. The used data sets are courtesy of Univ.-Klinik Innsbruck, AKH Vienna, Univ.-Klinikum Freiburg, and NLM.

References

- [GBKG] GRIMM S., BRUCKNER S., KANITSAR A., GRÖLLER E.: A refined data addressing and processing scheme to accelerate volume ray-casting. Accepted for Publication in *Computer & Graphics*.
- [GWGS02] GUTHE S., WAND M., GONSER J., STRASSER W.: Interactive rendering of large volume data sets. In *IEEE Visualization* (2002), pp. 53–60.
- [HMSC00] HUANG J., MÜLLER K., SHAREEF N., CRAWFIS R.: Fast splats: optimized splatting on rectilinear grids. In *IEEE Visualization* (2000).
- [Kni00] KNITTEL G.: The Ultravis system. In *IEEE Symposium on Volume visualization* (2000), pp. 71–79.
- [LL94] LACROUTE P., LEVOY M.: Fast volume rendering using a shear-warp factorization of the viewing transformation. In *Computer graphics and Interactive Techniques* (1994), pp. 451–458.
- [LY96] LAW A., YAGEL R.: Multi-frame trashless ray casting with advancing ray-front. In *Graphics Interfaces* (1996), pp. 70–77.
- [MJC02] MORA B., JESSEL J., CAUBET R.: A new object order ray-casting algorithm. In *IEEE Visualization* (2002), pp. 107–113.
- [MKW*02] MEISSNER M., KANUS U., WETEKAM G., HIRCHE J., EHLERT A., STRASSER W., DOGGETT M., FORTHMANN P., PROKSA R.: Vizard II, a reconfigurable interactive volume rendering system. In *Eurographics Workshop on Graphics Hardware* (2002), pp. 137–146.
- [MMC99] MÜLLER K., MÖLLER T., CRAWFIS R.: Splatting without the blur. In *IEEE Visualization* (1999), pp. 363–370.
- [MMMY97] MÖLLER T., MACHIRAJU R., MÜLLER K., YAGEL R.: A comparison of normal estimation schemes. In *IEEE Visualization* (1997), pp. 19–26.
- [MSHC99] MÜLLER K., SHAREEF N., HUANG J., CRAWFIS R.: Highquality splatting on rectilinear grids with efficient culling of occluded voxels. In *IEEE Transactions on Visualization and Computer Graphics* (1999), pp. 116–134.
- [NCV*02] NEUMANN L., CSEBFALVI B., VIOLA I., MLEJNEK M., GRÖLLER E.: Feature-preserving volume filtering. In *Symposium on Visualization* (2002), pp. 105–114.
- [PHK*99] PFISTER H., HARDENBERGH J., KNITTEL J., LAUER H., SEILER L.: The VolumePro real-time ray-casting system. In *SIGGRAPH, Computer Graphics* (1999), pp. 251–260.
- [PSL*98] PARKER S., SHIRLEY P., LIVNAT Y., HANSEN C., SLOAN P.: Interactive ray tracing for isosurface rendering. In *IEEE Visualization* (1998), pp. 233–238.
- [SFH97] SRINIVASAN R., FANG S., HUANG S.: Volume rendering by template-based octree projection. In *Eurographics, Visualization in Scientific Computing* (1997).
- [WE98] WESTERMANN R., ERTL T.: Efficiently using graphics hardware in volume rendering applications. In *SIGGRAPH* (1998), pp. 169–177.
- [Wes90] WESTOVER L.: Footprint evaluation for volume rendering. In *Computer Graphics and Interactive Techniques* (1990), pp. 367–376.
- [WG92] WILHELMS J., GELDER A. V.: Octrees for faster isosurface generation. In *Transactions on Graphics* (1992), pp. 201 – 227.