Hybrid Data Visualization Based On Depth Complexity Histogram Analysis

Stefan Lindholm, Martin Falk, Erik Sundén, Alexander Bock, Anders Ynnerman and Timo Ropinski

Linköping University Post Print

N.B.: When citing this work, cite the original article.

Original Publication:

Stefan Lindholm, Martin Falk, Erik Sundén, Alexander Bock, Anders Ynnerman and Timo Ropinski, Hybrid Data Visualization Based On Depth Complexity Histogram Analysis, 2014, Computer graphics forum (Print), (34), 1, 74-85.

<http://dx.doi.org/10.1111/cgf.12460>

Copyright: Wiley: 12 months

<http://eu.wiley.com/WileyCDA/>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-110238>

Hybrid Data Visualization Based On Depth Complexity Histogram Analysis

S. Lindholm, M. Falk, E. Sundén, A. Bock, A. Ynnerman, and T. Ropinski

Department of Science and Technology, Linköping University
{stefan.lindholm, martin.falk, erik.sunden, anders.ynnerman, timo.ropinski}@liu.se

Abstract

In many cases, only the combination of geometric and volumetric data sets is able to describe a single phenomenon under observation when visualizing large and complex data. When semi-transparent geometry is present, correct rendering results require sorting of transparent structures. Additional complexity is introduced as the contributions from volumetric data have to be partitioned according to the geometric objects in the scene. The A-buffer, an enhanced framebuffer with additional per-pixel information, has previously been introduced to deal with the complexity caused by transparent objects. In this paper, we present an optimized rendering algorithm for hybrid volume-geometry data based on the A-buffer concept. We propose two novel components for modern GPUs that tailor memory utilization to the depth complexity of individual pixels. The proposed components are compatible with modern A-buffer implementations and yield performance gains of up to eight times compared to existing approaches through reduced allocation and reuse of fast cache memory. We demonstrate the applicability of our approach and its performance with several examples from molecular biology, space weather, and medical visualization containing both, volumetric data and geometric structures.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism— I.3.8 [Computer Graphics]: Applications—

1. Introduction

With the widespread use of imaging technologies in data intensive research fields, visualization becomes an enabling technology that supports the exploration of acquired data sets. A recent trend is the multitude of data that arise not only in the form of multimodal volumetric data sets but also from geometric representations, which may be derived from the data or are acquired in different ways. Consequently, a challenge in visualization research is the efficiency and effectiveness at which this hybrid data can be rendered in a single setting. In many cases the integrated geometry is of similar or higher complexity than the volumetric data with which it should be combined. While modern volume rendering supports the inspection of otherwise dense data sets by facilitating semi-transparency, additional challenges arise when multiple volumes are combined with geometric data. Recent work suggests that geometric representations also benefit from semi-transparent properties [GRT13]. Nevertheless, current state-of-the-art algorithms cannot always deliver in-

teractive exploration for complex and semi-transparent hybrid data. Interactivity is, however, an essential feature when exploring scientific data sets.

In this paper, we improve upon existing techniques to facilitate efficient visualization of hybrid data sets as they occur in today's imaging-dependent areas of science. The fundamental challenge, arising when blending multiple semi-transparent sources into a single image, is to ensure that the elements are blended in view-dependent front-to-back or back-to-front order. When fusing multiple volumes or including mesh geometry into the visualization pipeline, the sorting must be ensured along each viewing ray. This results in a performance bottleneck as well as in increased memory usage. As these shortcomings hinder interactive exploration of many hybrid data sets, we have analyzed several of these data sets from different real-world scenarios with the goal to develop an optimized rendering algorithm enabling interactive exploration. As a tool in the analysis we have used Depth Complexity Histograms (DCH) that corre-

spond to the observed depth complexity across all pixels in the rendered image for a given scene and camera setting. The DCHs have been used to identify trends in the global distribution for scenes with hybrid data and to determine how these trends relate to rendering complexity. Based on the similarities of the occurring distributions, we propose two novel components which we combine with existing techniques to form an optimized rendering algorithm for hybrid data. The first component minimizes the allocated size in fast cache memory by adjusting the allocation to pixel depth complexity. The second component partitions the depth sorting similarly to depth peeling and recycles a smaller amount of allocated memory. As a result, computational throughput is increased and the maximum supported depth complexity is less dependent on the amount of available local GPU memory. The components provide between three and eightfold performance increase compared to existing algorithms. These qualities enable the interactive visualization of large and complex hybrid data sets at interactive frame rates, which is one of the key ingredients for enabling scientific discoveries in data-intense sciences.

2. Related Work

Data fusion. One aspect of hybrid data rendering relates to fusion of multiple volumetric sources. Several publications have investigated how samples from multiple sources should be blended, including different levels of intermixing [CS99]. A recent comparison of different intermixing schemes was presented by Schubert and Scholl [SS11].

Rendering multiple volumes with arbitrary placement, orientation, and resolution is often associated with computational costs for scene partitioning. To limit the per frame rendering cost, object space is often separated into convex regions homogeneously occupied by a fixed number of volumes. The convex regions are then sorted, either using specialized data structures [GBKG04, LLHY09, LF09] or using depth peeling techniques [PHF07, RBE08]. Many of these approaches also utilize shader instantiation, similar to one of the components presented in this paper, but do not include solutions for inclusion of semi-transparent geometry.

Fusing opaque geometry with volumetric data is straightforward in most volume rendering pipelines and does not require a full scene partition [EHK*06]. Semi-transparent geometry, on the other hand, is more computationally demanding as the resulting object space partitioning is costly, particularly if the geometry is not closed or has concave features. Image space techniques have been used in such situations as these techniques are generally well adapted to inclusion of generic, semi-transparent geometry [BBP+HR08, KGB*09]. With these methods, entry and exit points for homogeneously occupied ray segments are extracted from the volume proxy geometry and scene partitioning is performed on a per-ray basis during rasterization. The advantages and disadvantages of the available image space approaches are di-

rectly related to which transparency technique they employ, i.e. depth peeling and A-buffers, which are discussed below.

Transparency rendering. To ensure correct blending order of multiple semi-transparent samples, Order Independent Transparency (OIT) algorithms have been an ongoing research topic for the last thirty years [MCTB11]. Two of the most widespread approaches are depth peeling and A-buffers. Our algorithm follows the principles and global memory management of the A-buffer, which yields superior performance for scenes with high depth complexity [YHGT10, KKP*13].

The A-buffer is a leading OIT solution based on the temporary storage and sorting of intermediate pixel fragments. The concept was first introduced by Carpenter [Car84]. Subsequent works improved the concept [EP90, BCL*07, MB07, KWBG13, VF14], including hardware adaptations that were generally limited to a comparatively low number of samples or utilized pre-sorting of primitives. Further improvements in graphics hardware led to A-buffer variants with support for higher depth complexities including paged variations of the algorithm [KGB*09, Cra10], Per-Pixel Linked Lists (PPLL) [YHGT10], and an alternative called Dynamic Fragment Buffers (DFB) or s-buffer [MCTB12, VF12]. Both DFB and PPLL approaches will be described later in detail along with the straightforward Fixed Fragment Buffer (FFB) implementation [Cra10]. However, it should be noted here that the modern A-buffer literature focuses only on management of global GPU memory. Our work builds upon these methods by improving the management of local GPU caches and hence improves performance for a wide range of A-buffer implementations.

Depth peeling was first introduced by Mammen [Mam89] and has evolved to employ multi-directional and bucket-approximations [Eve01, CMM08, BM08, LHLW09]. Although our main algorithm is not based on depth peeling, the technique arguably remains a close competitor to A-buffers among OIT solutions and remains an active area of research [VF13]. In addition, one of the components proposed in this paper shares high resemblance with the original depth peeling approach.

3. Visualization of Scenes with Non-uniform Depth Complexity

In this paper, we have chosen four representative scenes of different fields including molecular science (*protein*), space weather simulations (*space*), neurosurgical treatment planning (*medical*), and computational fluid dynamics (*cfid*). Characteristic for the selected scenes is that the major computational costs arise from the task of ordering all contributions in terms of visibility. In modern OIT solutions, these costs are directly related to the observed depth complexity in the image, i.e., the number of overlapping contributions per pixel. At the same time, most scenes often feature non-

Figure 1: Visualization of a protein (*human carbonic anhydrase II*) combining hybrid data from different sources. (a) final rendering of hybrid visualization. (b) geometric protein representations. (c) volumetric data sources depicting electrostatic potentials. (d) image-space depth complexity (re-scaled for representational purposes). (e) depth complexity histogram (DCH, log scale). Our improved A-buffer algorithm is capable of rendering the entire scene at 32 fps—a performance increase of 5.9 times compared to prevalent techniques.

uniform complexity distributions in image space with clusters of high depth complexity surrounded by regions of lower complexity.

The first hybrid data scene is shown in Figure 1 and depicts the protein *human carbonic anhydrase II* (PDB-ID: 2CBA). The scene contains protein structures stored in polygonal meshes and volumetric data sources representing electron charge densities. Such protein visualizations are commonly used to examine possible docking positions between the protein and other molecules [SdG10]. To analyze the complexity of such scenes we thus utilize depth complexity histograms.

3.1. Depth Complexity Histograms

As an example for the Depth Complexity Histogram (DCH), we use the protein scene in Figure 1, where the DCH for the shown camera setting is depicted in (e). The DCH is computed by first rendering a depth complexity image as shown in (d) (note that the image shown has been normalized and inverted for presentational purposes). When rendering the complexity image, both shading and blending are disabled. All geometries and volume bounding boxes are rendered into an integer image aligned with the current viewport. The pixels of this image act as atomic counters and are incremented by the incoming fragments. Each pixel of the complexity im-

age thus holds the number of fragments. The DCH is then computed as a standard histogram over this integer image. Each bin in a DCH is thereby associated with a complexity range and holds the number of pixels whose depth complexities fall within that range.

3.2. Visualization Challenges for Data Fusion

During the analysis of our four scenes, we made the observation that many pixels feature a low depth complexity and only a few pixels feature the highest depth complexities. This results in rapidly decreasing DCHs where the majority of pixels have a complexity that is only a fraction of the global maximum of the scene. In our experience, these findings are also applicable to many other scenes.

Again, we use the protein scene in Figure 1 as an example. From a depth complexity perspective, the protein scene is fairly well behaved. The maximum depth complexity is limited and the points with high depth complexity are evenly distributed across the entire scene. The DCH, however, decreases rapidly. Note that the DCH in Figure 1(e) is plotted logarithmically. The dominance of even-numbered depth complexities is a result of closed geometries. The majority of all pixels (68%) have a depth complexity of 8 or less, whereas only very few pixels (6.6%) have a complex-

Figure 2: Performance and complexity over time for a pre-defined camera path around the protein data set of Figure 1. Rendering times are shown as lines (first y-axis) and DCHs as bar charts (second y-axis). The camera dollies in and out while the scene composition is changed every 100th step.

ity larger than 17 with a maximum of 46 for this specific camera setting.

DCHs are not only dependent on scene content but also on the camera settings. Figure 2 illustrates how the DCH changes for the protein scene for an animated camera path. The time-dependent DCHs are shown as vertical bar-plots of screen coverage in percent (linear scale, right y-axis). Complexities have been binned according to the legend in the lower right and the DCH is plotted for every tenth time step. The camera path is illustrated with five snapshots taken at time points 50, 150, 250, 350, and 450 respectively (top row of images). Additionally, the figure also shows performance (red and blue line plots, left y-axis) which will be discussed in later sections. In addition to camera movement, the scene composition is changed every 100th time-point. At the first marker (orange), the geometry of the stick representation is removed (cf. Figure 1(b)) reducing high-complexity pixels. At the second marker (green), both volume data sources are removed (cf. Figure 1(c)) reducing low-complexity pixels. From the plots, we can see that the DCH changes over time and that the rendering performance depend on these changes. We can also see that the DCH remains non-uniform in nature even with the camera very close to the protein. This opens

Figure 3: Overview of the A-buffer rendering approach on GPU hardware. Previous work has primarily focused on global memory management during the Fill Step. Our proposed algorithm is hence aimed at local GPU cache memory. Memory related steps are highlighted in blue.

up the possibility to save resources if the memory allocation could be adapted to the observed complexity on a per pixel basis. In the next section we provide a technical description of how current state-of-the-art OIT techniques adapt to non-uniform scene complexities.

4. Algorithms for Semi-Transparent Data Fusion

To construct an optimized algorithm we exploit the trend of non-uniform DCHs observed earlier. Three state-of-the-art implementations of the A-buffer concept serve as a reference. This section details the differences of those algorithms in the dynamic management of global GPU memory. In addition, we show how to further optimize the algorithms with respect to their use of local caches on modern graphics hardware.

The principle of the A-buffer is to capture and store a list of rasterized fragments on a per-pixel basis (called *Fragment List*). This involves three sequential steps as illustrated in Figure 3: the Clear Step, the Fill Step, and the Resolve Step. Both the Fill Step and Resolve Step contain sub-steps associated with memory management (highlighted in blue). This includes global memory management mainly during the Fill Step and local cache management in the Resolve Step. Fragment data structure and further considerations on source types will be described in Section 6.

Managing Global GPU Memory

Three prevalent approaches for managing the global A-buffer memory during the Fill Step are used. These are Fixed Fragment Buffers (FFB) [Cra10], Dynamic Fragment Buffers (DFB) [MCTB12, VF12], and Per-Pixel Linked Lists (PPLL) [KGB*09, YHGT10, Cra10].

In Figure 4, the three approaches are illustrated side-by-side. In each approach, every pixel is associated with a fragment counter and an implicit or explicit pointer to the stored

Figure 4: Comparison of three existing A-buffer implementations for managing global GPU memory. Both, the DFB and PPLL approaches adapt well to scenes with non-uniform depth complexities by allowing dynamic memory utilization.

Fragment List. For the FFB, the global memory pool is statically partitioned and the pointer is implicitly given by the screen coordinates of the fragments. The global memory pool of the DFB is dynamically partitioned each frame and the explicit pointer holds a base offset into the pool. For the PPLL, the pointer denotes the list anchor for a linked list of fragments spread throughout the pool. Both, FFB and DFB ensure a contiguous memory layout for each individual pixel. The Fragment Lists are generally stored out-of-order in the comparatively slow but large global memory of the GPU.

DFB and PPLL are both dynamic algorithms and adapt the utilization of global memory well to non-uniform depth complexity distributions. Thus, consumption of global memory is almost optimal even for rapidly decreasing DCHs.

Managing Local GPU Caches

It is common practice to avoid the higher latency of global memory by copying the unsorted Fragment List to a local cache before sorting. An important aspect of such cache utilization is that the memory layout of modern GPUs typically exposes a limited shared cache to a group of cores. Additionally, modern GPUs hide global memory latency by hot swapping groups of threads in a manner similar to pipelining on a CPU [NV11]. Under ideal circumstances, the number of active threads can be eight times higher than the number of physical cores. Per thread allocations therefore need to be small enough such that all active threads designated to a group of cores can have their arrays allocated simultaneously. It is thus important to minimize the allocation of cache memory to maximize performance. Current trends also indicate that the number of cores increases faster than the available shared cache size, potentially escalating this problem in the future.

The prevalent approach described in the literature for managing local GPU caches is to allocate a fixed sized array, of size N per thread [Cra10, MCTB12]. This has two immediate consequences. First, the maximum depth complexity handled by the approach is limited to N , since all fragments beyond this number are discarded. Second, larger array sizes significantly reduce the number of active threads

due to cache overflow and reduced hot swapping. A number often reported in the literature is $N = 64$. While the literature includes strategies, such as DFB and PPLL, for managing the global memory, there still remains room for improvement in managing local caches.

5. Improved Dynamic Depth Complexity Management

We propose two novel A-buffer components for A-buffer based algorithms based on the observed nature of the DCHs. Both components are designed to improve the management of local GPU caches. For clarity, we will use the term Local Array to indicate the fast (local) memory allocated for the sorting of fragments.

The first component, illustrated in the center of Figure 5, is called per-pixel Array Optimization (ppAO) and ensures that all pixels are evaluated without allocating excessively large Local Arrays. The component segments image space into sections of similar depth complexity and separate shaders, compiled with different local arrays sizes, are triggered for each segment. The second component, illustrated to the right in Figure 5, is part of the Resolve Step and corresponds to a novel sorting procedure called per-pixel Depth Peeling (ppDP). The component breaks the sorting into smaller pieces by not loading all fragments at once and can thereby recycle memory.

A key factor of both components is that they are designed to be used in combination with the current state-of-the-art management of global memory described in the previous section. The full rendering algorithm, optimized for geometry intensive fusion scenes, is thus formed by combining the components proposed here with one of the existing solutions for global management, such as DFB or PPLL.

5.1. Dynamic Resource Management Using per-pixel Array Optimization

The objective of the proposed ppAO component is to limit the amount of unused memory in the local cache by performing the Resolve Step with Local Arrays that better correspond to the depth complexity of individual pixels. Due

Figure 5: Comparison of approaches managing local GPU caches: static management (left) and our two proposed components ppAO (center) and ppDP (right). The unsorted input Fragment List (top) corresponds to the output of the Fill Step and resides in global memory. The array is sorted in local caches during the Resolve Step using predefined array sizes. Two pixels are exemplified, a heavy pixel (blue) and a light pixel (orange). The ppAO component significantly lowers the amount of unused memory (shown in gray), while the ppDP component lowers the amount of both unused and used memory by recycling it.

to restrictions in GPU cache allocations, the optimization of per-pixel array sizes requires multiple shader programs to be instantiated for different Local Array sizes. Rather than using a unique shader for every possible Local Array size we use a smaller set of pre-defined array sizes. This leads to a clustering of pixels with similar depth complexity into segments. All pixels in a segment are then resolved in a batch by a single shader. The full image is then processed in as many batches as there are segments. Note that the clustering is independent of the actual pixel locations and that each segment not necessarily corresponds to a focused region in the image. The ppAO is positioned after the Fill Step in the A-buffer algorithm (cf. Figure 3) making the ppAO responsible for managing and calling the Resolve Step on the respective complexity segments. The ppAO algorithm is outlined below.

```
// Fill Step
forall the pixels of the framebuffer do
  | count depth complexity;
end
// Resolve Step, loop over all complexity segments
foreach  $i \in \{8, 16, \dots, \text{max}\}$  do
  activate shader program for complexity segment  $i$ ;
  forall the unresolved pixels of the framebuffer do
    // mask pixel if not part of segment class  $i$ 
    if depth complexity( $p$ )  $\leq i$  then
      | resolve segment with Local Array of size  $i$ ;
      | // mark pixel as resolved
    end
  end
end
end
```

The instantiation procedure of the shader programs is straightforward as the individual instances only differ by a single number (the predetermined buffer size). It is also suf-

ficient to instantiate only a small set of shaders due to the decreasing power curve observed in most DCHs. By analyzing the depth complexity of various scenes, we found that complexity segments of sizes 8, 16, 32, $\dots$, N give good results with respect to performance and memory consumption.

To perform ppAO, we need to know the depth complexities of all pixels. For this purpose, we utilize a separate integer buffer the size of the framebuffer (one atomic counter per pixel) to count the number of fragments per pixel. The counting can be combined with the scene rasterization in the Fill Step and, hence, no additional rendering passes are required. Once the depth complexities are known, the different complexity segments may be processed. This is done by looping over all pre-defined sizes of complexity segments and activating the shader program with the respective Local Array size. The clustering of pixels into complexity classes is handled implicitly during runtime by a masking step where all pixels that are not part of the current segment are discarded. Note that the graphics pipeline does not need to be flushed between triggering successive segments and that multiple segments may be evaluated in parallel (even if they are triggered sequentially). Figure 5, center, depicts the utilization of our approach for pixels exhibiting high and low depth complexity.

5.2. Preventing Over-sized Local Arrays Using per-pixel Depth Peeling

If the number of fragments inside a Fragment List exceeds the buffer size of the Local Array information is usually lost and the final rendering result might feature artifacts. One impractical solution is to increase the size of the Local Arrays. However, large parts of the pre-allocated memory will never be used since typically only a few pixels exhibit a high depth complexity (cf. Section 3.1 and Figure 5, left). To overcome

this particular problem, we propose a variation of depth peeling on a per-pixel basis to provide a simple way to correctly deal with overflowing Local Arrays. Thus, the maximum supported depth complexity is no longer limited by the size of local memory but instead limited only by the amount of available global memory. In Figure 5, right, a Local Array size of 8 is used to illustrate our approach for two depth complexities.

The ppDP algorithm is outlined below and replaces the Resolve Step for individual pixels in Figure 3.

```

input : list of fragment data per pixel
 $n \leftarrow \text{sizeof}(\text{Local Array});$ 
create buffer with  $n$  elements;
set min depth to 0;
initialize current state;
// Resolve Step
repeat
  foreach fragment  $f \in \text{fragment list}$  do
    if  $\text{depth}(f) \leq \text{min depth}$  then
      | skip  $f$  in current pass;
    else
      | insert  $f$  into buffer with insertion sort;
      | // store at most  $n$  elements
      | // drop last element on overflow
    end
  end
  resolve buffer considering current state;
  store min depth and state of last buffer element;
  clear buffer;
until all fragments  $f \in \text{fragment list}$  are resolved;

```

The ppDP algorithm follows the approach of bucket depth peeling where a larger list is sorted and resolved as a set of smaller sequences but does so fully on the GPU between global and local memory without the need to re-render geometry. We split the Resolve Step into several subpasses which allows for the Local Array to be recycled and its size to be reduced. With a Local Array of size n , the process sorts $n - 1$ elements in the Fragment List per subpass. While looping over the contents of the Fragment List, the n entries with the smallest depth values are chosen and stored in the buffer by using insertion sort. If a fragment was already processed it will be discarded in the current pass. After filling the buffer its contents are resolved. To ensure consistency between the individual resolve passes, we temporarily store the current state of the rendering, including volume occupancy and ray position. The current state and minimal depth are updated and the local buffer is emptied for the next pass. This approach does not require additional rendering passes but does require the entire Fragment List to be read multiple times from global GPU memory.

Sorting the full array thus takes a maximum of $\lceil D/(n - 1) \rceil$ passes for a pixel with depth complexity D . The loop may be terminated sooner, e.g. in case of early-ray ter-

mination, and further read/write intensive sorting operations are avoided. Note that the size of the Local Array is constant. We observed an optimal Local Array size of 8 for our ppDP approach across all test setups with respect to performance.

Note that our ppDP algorithm may exhibit z-fighting issues similar to regular depth peeling. The issue arises when multiple fragments of a single pixel have identical depth and only a subset of these fragments gets included in one loop iteration. Since only the minimal depth is used to distinguish processed fragments from fragments yet to be resolved, the solution is not unique. This also complicates the decision in the subsequent iterations regarding which fragments should be discarded. So far in our work we have not experienced any artifacts from this kind of z-fighting. Potential solutions exist in the form of additional per-fragment flags which could be used to assess process status at the cost of additional read/write operations. A comprehensive discussion on the phenomenon and its solutions can be found in the work of Vasilakis and Fudos [VF13].

6. Fused Rendering of Hybrid Data Sources

So far we have discussed the memory management of the A-buffer involved in the construction of the sorted Fragment List (Figure 3, blue highlights). In this section, we will focus on the A-buffer evaluation and the rendering of the scene (Figure 3, orange highlight) and implementation details including data structures.

6.1. Fragment List Creation and Evaluation

The following data structure is used for all fragments generated in the Fill Step as entries of the Fragment List: z (float32), id (int32), $color$ (vec4 float16). The structure has a total size of 16 bytes and the same data structure is used for both geometric and volumetric contributions.

In the Fill Step, ABuffer_frag entries are computed and stored in the Fragment List. Geometric sources are shaded during this step and the computed color is stored explicitly in the $color$ component. For volumetric sources, only the associated proxy geometry is rendered during the Fill Step and the $color$ component is left uninitialized. The z and id components are treated identically for all source types and respectively hold the depth value in screen space coordinates and a unique source identifier. Bit operations are used on the id component to store both source type as well as a unique identifier. Sources may be rendered independently by different shaders but the produced fragments (instances of ABuffer_frag) are all inserted into the same global storage.

During the final stage of the Resolve Step, the scene content is evaluated and blended into the framebuffer. At this point, all entries in the Fragment List can be interpreted as intersection points along the view rays. Ray casting of the scene is performed in world space by sequentially looping

over all entries in the Fragment List. Geometry fragments are blended directly to the buffer while volume rendering is performed for the interleaved ray segments. Volume occupancy is tracked through a bitmask which is updated every time a fragment of a volume proxy geometry is encountered. For more information on fragment list evaluation, particularly the use of bitmasks to track volume occupancy, we direct the reader to Brecheisen *et al.* [BBPtHR08]. Once the occupancy and entry and exit points are known, the problem is reduced to the problem of multi-volume rendering as discussed in Section 2.

6.2. Implementation Details

Our implementation supports global memory management using FFB, DFB, and PPLL as described in Section 4. The FFB and PPLL implementations are derived from the work by Crassin [Cra10]. The code for the DFB implementation was provided by Maule *et al.* [MCTB12] and was slightly changed to fit our needs. We use C++, OpenGL, and GLSL for our framework. The only exception is the scan step of the DFB implementation which is performed in CUDA using the CUDA Thrust library as explained by Maule *et al.*

For management of local cache memory we utilize our proposed components: ppAO and ppDP (cf. Section 5). The ppAO component is implemented using an 8bit stencil buffer for masking the different complexity segments to be triggered by separate shaders. The Fill Step is enclosed in a loop on the CPU, where different segments are triggered by modifying the OpenGL stencil function. Note that using an 8bit buffer does not limit the depth complexity to 255. Higher complexities are still possible but they will all be associated with the same segment. If larger segment limits are needed, the stencil buffer can easily be replaced with a texture of higher precision. The ppDP approach simply replaces the final steps of the Resolve Step and the loop is implemented directly in GLSL.

Local Arrays are always allocated as local buffers with predefined sizes inside the shaders and are, thus, not assigned explicitly to specific memory. Thread scheduling and memory assignment is therefore up to the discretion of the driver. The entire algorithm, with global and local optimizations, requires the implementation of three shaders namely clear, fill, resolve plus additional clones of the resolve shader containing different buffer sizes for the ppAO.

7. Results

We tested our proposed A-buffer components with four real-world cases from different fields. The performance was measured for a viewport size of 1024×768 on four different NVIDIA GeForce GTX GPUs: 560, 580, 670, and Titan (1 GiB, 1.5 GiB, 2 GiB, and 6 GiB VRAM respectively). Individual performance for the 580 and Titan are included in

the paper (representing Fermi and Kepler architectures respectively). Performance for 560 and 670 are available as supplementary material. Reported averages were computed across all GPUs. Scene configurations will be described next in Section 7.1 followed by a presentation of performance in Section 7.2.

7.1. Scenario Descriptions

Figure 1, protein: The data corresponds to the protein *human carbonic anhydrase II*. The scene consists of four separate data sources; two volumetric data sets describing the electrostatic potential calculated at different resolutions and extents, a geometric stick representation of the protein, and a geometric ribbon model of the protein. Both geometric representations are colored according to the secondary structure type. The three-dimensional structure and potential fields are often visualized to examine potential docking positions between the protein and other molecules [SdG10].

Figure 8, space: The space data set depicts results of a multi-variate simulation of a time-dependent 3D magnetohydrodynamics system of the heliosphere during a coronal mass ejection event. This particular system is currently used in space weather prediction [XOL04]. The scene consists of two volumes, showing the number of charged particles and the energy density, and three isosurfaces derived from the energy densities.

Figure 9, medical: The data, provided as the IEEE Visualization Contest in 2010, contains multiple medical imaging modalities as well as derived information sources for planning neurosurgical intervention [CK10]. The scene consists of four data sources; two volumetric data sets depicting T1-weighted MR images of head and brain, two sources of geometric information in a surface extraction of a tumor segmentation and about 1678 fiber tracts obtained with diffusion tensor imaging. The combined information from all data sources is used to plan the safest possible access path for an intervention.

Figure 10, cfd: The data results from a computational fluid dynamics simulation of blood flow in the carotid artery of a human subject. The scene consists of a single computed tomography image and two sets of geometric primitives. A sparse glyph representation and more than 7500 individual streamlines.

More details on the different data sources of the individual scenes are listed in Table 1. Depth complexities are available as single view DCHs in Figures 1 (protein), 8 (space), 9 (medical), and 10 (cfd), respectively.

7.2. Performance Comparison

We have investigated how our proposed components improve the performance for each of the three A-buffer algorithms for global memory management. For each algorithm,

Table 1: Data information for the four selected scenes, including the first frame of the sequence used for performance tests.

Scene	Data	Data size	1 st Frame
protein	VOL	$127 \times 127 \times 127$
	VOL	$71 \times 71 \times 71$	
	GEO	868k triangles	
	GEO	36k triangles	
space	VOL	$256 \times 256 \times 256$
	VOL	$256 \times 256 \times 256$	
	GEO	162k triangles	
	GEO	100k triangles	
medical	VOL	$415 \times 487 \times 176$
	VOL	$367 \times 395 \times 150$	
	GEO	363k triangles	
	GEO	1678 fibers	
cfd	VOL	$76 \times 49 \times 45$
	GEO	7.5k streamlines	
	GEO	≈ 100 arrows	

we measured the performance for different combinations. In Figure 6, the results are shown from top to bottom for the four distinct scenes—protein, space, medical, and cfd. For each of FFB, DFB, and PPLL, we have measured the performance of four configurations for local memory management; baseline (static buffers), ppAO, ppDP, and ppAO+ppDP. The largest Local Array size allocated for the scenes was set to 64 (protein), 32 (space), and 128 (neuro and flow). Allocating a Local Array size to hold all depth layers of the medical scene resulted in driver crashes, thus, only the first 128 layers were resolved for baseline and ppAO configurations (ppDP is capable of resolving higher complexities with smaller Local Arrays). Likewise, using FFB for global management also limits maximum complexity, requiring approximately 12.6 MiB per depth layer at 1024×768 resolution. Depth segments for ppAO were chosen as powers of two between 8 and the scene cap reported above. All performance results for ppDP were measured with a Local Array size set to 8. For each test, frame times were computed as the average over 30 seconds of rendering as the camera rotated around the scene at fixed distance. Early-ray termination was enabled for transparency values less than 0.02 during all benchmarks.

Due to a performance penalty associated with a CUDA/OpenGL context switch, we were unable to achieve results for the DFB on the same level as reported by Maule *et al.* [MCTB12]. The penalty manifested as a 50ms–200ms delay associated with the scan step and was persistent across different GPUs and drivers and was also present in a min-

Figure 6: Performance comparison with our proposed ppAO and ppDP components. Columns represent the different A-buffer implementations in combination with our components (none, ppAO, ppDP, ppAO and ppDP combined).

imal stand-alone DFB implementation. The results for the DFB approach reported in Figure 6 include this penalty.

Supported by the performance graphs, we can say that both ppAO and ppDP are capable of providing significant performance gains for a large variety of configurations. The performance gain is threefold for ppAO and about eightfold ppDP when averaged over all scenes and GPUs. For scenes with high depth complexities or homogeneous depth complexity distributions, the improvement obtained with ppAO is reduced since the processing of the heavy pixels becomes a bottleneck. In such scenarios, ppDP maintains a significant speedup, particularly when early-ray termination is enabled.

In addition to the four real-world scenarios, we created a synthetic worst-case scene with screen-filling quads. The scene consists of 64 quads ($\alpha = 0.05$) aligned orthogonal to the viewing direction where their order is randomized in depth. The quads are scaled so that they all cover the entire viewport in the middle of the animation. Thus, all pixels feature a near-maximum depth complexity. Results are shown in Figure 7 and confirm that the performance of ppAO (as expected) approaches that of statically implemented buffers while ppDP maintains a 3.6 times performance increase.

Figure 7: Performance for a synthetic data set. Rendering times are shown as lines (first y-axis) and DCHs as bar charts (second y-axis).

8. Conclusions and Future Work

In this paper, we address the challenge of interactive exploration of complex hybrid data. Two adaptive data handling components are proposed to be combined with prevalent A-buffer techniques to form a final rendering algorithm. Of the two, ppDP is better equipped to deal with very high complexities but does so at the cost of potential artifacts from z-fighting while ppAO performs well for moderate complexities with exact results. Optionally, the components can be combined to achieve support for high depth complexities while guaranteeing a minimum number of correctly blended samples.

An interesting consequence of our work is the potential alleviation of the traditional trade-off between performance and depth-support. We believe that this aspect makes our work valuable also in other areas of visualization or computer graphics that employ A-buffer based techniques. For example, it may now be possible to encode fiber tracking uncertainties in the alpha channel of the visualized fibers without eliminating the possibility to also visualize a co-registered volume.

Acknowledgments

Space simulation data have been provided by CCMC at Goddard Space Flight Center, NASA. This work was supported in part by the Swedish Research Council, VR grant 2011-5816 and the Linnaeus Environment CADICS, and the Swedish e-Science Research Centre (SeRC).

References

- [BBPthR08] BRECHEISEN R., BARTROLI A. V., PLATEL B., TER HAAR ROMENY B. M.: Flexible GPU-based multi-volume ray-casting. In *Vision, Modelling and Visualization (VMV)* (Konstanz, Germany, 2008), pp. 303–312. 2, 8
- [BCL*07] BAVOIL L., CALLAHAN S. P., LEFOHN A., COMBA J. A. L. D., SILVA C. T.: Multi-fragment effects on the GPU using the k -buffer. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Seattle, USA, 2007), pp. 97–104. doi:10.1145/1230100.1230117. 2
- [BM08] BAVOIL L., MYERS K.: *Order-Independent Transparency with Dual Depth Peeling*. Tech. rep., NVIDIA Developer SDK 10, 2008. URL: <http://developer.download.nvidia.com>. 2
- [Car84] CARPENTER L.: The A-buffer, an antialiased hidden surface method. In *ACM SIGGRAPH Computer Graphics and Interactive Techniques* (Minneapolis, USA, 1984), pp. 103–108. doi:10.1145/800031.808585. 2
- [CK10] CHOURASIA A., KLEIN J.: IEEE visualization contest 2010. Online resource, Oct 2010. URL: <http://sciviscontest.ieeevis.org/2010/>. 8
- [CMM08] CARR N., MÉCH R., MILLER G.: Coherent layer peeling for transparent high-depth-complexity scenes. In *ACM SIGGRAPH/EUROGRAPHICS Symposium on Graphics Hardware* (Sarajevo, Bosnia-Herzegovina, 2008), pp. 33–40. doi:10.2312/EGGH/EGGH08/033-040. 2
- [Cra10] CRASSIN C.: OpenGL 4.0+ ABuffer v2.0: Linked lists of fragment pages. Personal Blog, July 2010. Open source OIT implementation. URL: <http://blog.icare3d.org/2010/07/>. 2, 4, 5, 8
- [CS99] CAI W., SAKAS G.: Data intermixing and multi-volume rendering. *Computer Graphics Forum* 18, 3 (1999), 359–368. doi:10.1111/1467-8659.00356. 2
- [EHK*06] ENGEL K., HADWIGER M., KNISS J. M., REZK-SALAMA C., WEISKOPF D.: *Real-Time Volume Graphics*. A. K. Peters, Ltd., Natick, MA, USA, 2006. 2
- [EP90] EBERT D. S., PARENT R. E.: Rendering and animation of gaseous phenomena by combining fast volume and scanline a-buffer techniques. In *Proceedings of SIGGRAPH* (Dallas, USA, 1990), ACM, pp. 357–366. doi:10.1145/97880.97918.2
- [Eve01] EVERITT C.: *Interactive Order-Independent Transparency*. Tech. rep., NVIDIA White Paper, 2001. URL: <https://developer.nvidia.com>. 2
- [GBKG04] GRIMM S., BRUCKNER S., KANITSAR A., GRÖLLER E.: Flexible direct multi-volume rendering in interactive scenes. In *Vision, Modelling and Visualization (VMV)* (Stanford, USA, Oct. 2004), pp. 386–379. 2
- [GRT13] GÜNTHER T., RÖSSL C., THEISEL H.: Opacity optimization for 3D line fields. *ACM Transactions on Graphics* 32, 4 (July 2013), 120:1–120:8. doi:10.1145/2461912.2461930. 1
- [KGB*09] KAINZ B., GRABNER M., BORNIK A., HAUSWIESNER S., MUEHL J., SCHMALSTIEG D.: Ray-casting of multiple volumetric datasets with polyhedral boundaries on manycore GPUs. *ACM Transactions on Graphics* 28, 5 (Dec. 2009), 152:1–152:9. doi:10.1145/1618452.1618498. 2, 4
- [KKP*13] KAUKER D., KRONE M., PANAGIOTIDIS A., REINA G., ERTL T.: Rendering molecular surfaces using order-independent transparency. In *Eurographics Symposium on Parallel Graphics and Visualization (EGPGV)* (Girona, Spain, 2013), vol. 13, pp. 33–40. doi:10.2312/EGPGV/EGPGV13/033-040. 2

Figure 8: Visualization of solar mass ejection simulation data used for predicting space weather. (a) final rendering of hybrid visualization. (b) three geometric isosurfaces. (c) volumetric data sources. (d) image-space depth complexity (re-scaled for representational purposes). (e) depth complexity histogram (DCH, log scale). Our improved A-buffer algorithm achieves up to three times the performance compared to existing approaches.

Figure 9: Fused visualization of hybrid data for the purpose of neurosurgical planning. The scene is particularly complex in areas where the DTI fibers converge toward two main bundles. (a) full scene, (b) zoomed view, and (c) the depth complexity histogram (DCH, log scale) of (a).

Figure 10: Fused visualization of hybrid data used to assess blood flow in the human carotid artery. The rendering uses transparency to reduce the focus on individual streamlines. (a) full scene, (b) zoomed view, and (c) the depth complexity histogram (DCH, log scale) of (a).

- [KWBG13] KERZNER E., WYMAN C., BUTLER L., GRIBBLE C.: Toward efficient and accurate order-independent transparency. In *ACM SIGGRAPH 2013 Posters* (Anaheim, USA, 2013), SIGGRAPH '13, ACM, pp. 109:1–109:1. doi:10.1145/2503385.2503504. 2
- [LF09] LUX C., FRÖHLICH B.: GPU-based ray-casting of multiple multi-resolution volume datasets. In *International Symposium on Advances in Visual Computing: Part II* (Las Vegas, USA, 2009), Springer-Verlag, pp. 104–116. doi:10.1007/978-3-642-10520-3_10. 2
- [LHLW09] LIU F., HUANG M.-C., LIU X.-H., WU E.-H.: Efficient depth peeling via bucket sort. In *Conference on High Performance Graphics* (New Orleans, USA, 2009), pp. 51–57. doi:10.1145/1572769.1572779. 2
- [LLHY09] LINDHOLM S., LJUNG P., HADWIGER M., YNNERMAN A.: Fused multi-volume DVR using binary space partitioning. *Computer Graphics Forum* 28, 3 (2009), 847–854. doi:10.1111/j.1467-8659.2009.01465.x. 2
- [Mam89] MAMMEN A.: Transparency and antialiasing algorithms implemented with the virtual pixel maps technique. *IEEE Computer Graphics and Applications* 9, 4 (1989), 43–55. doi:10.1109/38.31463. 2
- [MB07] MYERS K., BAVOIL L.: *Stencil Routed K-Buffer*. Tech. rep., NVIDIA White Paper, 2007. URL: <http://developer.download.nvidia.com>. 2
- [MCTB11] MAULE M., COMBA J. L., TORCHELSEN R. P., BASTOS R.: A survey of raster-based transparency techniques. *Computers And Graphics* 35, 6 (2011), 1023–1034. doi:10.1016/j.cag.2011.07.006. 2
- [MCTB12] MAULE M., COMBA J., TORCHELSEN R., BASTOS R.: Memory-efficient order-independent transparency with dynamic fragment buffer. In *Conference on Graphics, Patterns and Images (SIBGRAPI)* (Ouro Preto, Brazil, 2012), pp. 134–141. doi:10.1109/SIBGRAPI.2012.27. 2, 4, 5, 8, 9
- [NVI11] NVIDIA: Performance optimization. Tutorial at Super Computing Conference 2011, 2011. NVIDIA online document. URL: <http://www.nvidia.com/object/sc11.html>. 5
- [PHF07] PLATE J., HOLTKAEMPER T., FROELICH B.: A flexible multi-volume shader framework for arbitrarily intersecting multi-resolution datasets. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1584–1591. doi:10.1109/TVCG.2007.70534. 2
- [RBE08] RÖSSLER F., BOTCHEN R., ERTL T.: Dynamic shader generation for GPU-based multi-volume ray casting. *IEEE Computer Graphics and Applications* 28, 5 (2008), 66–77. doi:10.1109/MCG.2008.96. 2
- [SdG10] SEELIGER D., DE GROOT B. L.: Ligand docking and binding site analysis with PyMOL and Autodock/Vina. *Journal of Computer-Aided Molecular Design* 24, 5 (2010), 417–422. doi:10.1007/s10822-010-9352-6. 3, 8
- [SS11] SCHUBERT N., SCHOLL I.: Comparing GPU-based multi-volume ray-casting techniques. *Computer Science-Research and Development* 26, 1-2 (2011), 39–50. doi:10.1007/s00450-010-0141-1. 2
- [VF12] VASILAKIS A., FUDOS I.: S-buffer: Sparsity-aware multi-fragment rendering. In *Eurographics (Short Papers)* (Cagliari, Italy, 2012), Andújar C., Puppo E., (Eds.), Eurographics Association, pp. 101–104. doi:10.2312/conf/EG2012/short/101-104. 2, 4
- [VF13] VASILAKIS A., FUDOS I.: Depth-fighting aware methods for multifragment rendering. *IEEE Transactions on Visualization and Computer Graphics* 19, 6 (2013), 967–977. doi:10.1109/TVCG.2012.300. 2, 7
- [VF14] VASILAKIS A. A., FUDOS I.: k^+ -buffer: Fragment synchronized k-buffer. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (San Francisco, USA, 2014), I3D '14, ACM, pp. 143–150. doi:10.1145/2556700.2556702. 2
- [XOL04] XIE H., OFMAN L., LAWRENCE G.: Cone model for halo CMEs: Application to space weather forecasting. *Journal of Geophysical Research: Space Physics* 109, A3 (2004). doi:10.1029/2003JA010226. 8
- [YHGT10] YANG J. C., HENSLEY J., GRÜN H., THIBIEROZ N.: Real-time concurrent linked list construction on the GPU. *Computer Graphics Forum* 29, 4 (2010), 1297–1304. doi:10.1111/j.1467-8659.2010.01725.x. 2, 4