

# TECHNICAL REPORT

## Progressive Volume Rendering of Large Unstructured Grids

*Steven P. Callahan, Louis Bavoil, Valerio Pascucci, and Cláudio T. Silva*

UUSCI-2006-019

Scientific Computing and Imaging Institute  
University of Utah  
Salt Lake City, UT 84112 USA

May 2, 2006

### **Abstract:**

We describe a new progressive technique that allows real-time rendering of extremely large tetrahedral meshes. Our approach uses a client-server architecture to incrementally stream portions of the mesh from a server to a client which refines the quality of the approximate rendering until it converges to a full quality rendering. The results of previous steps are re-used in each subsequent refinement, thus leading to an efficient rendering. Our novel approach keeps very little geometry on the client and works by refining a set of rendered images at each step. Our interactive representation of the dataset is efficient, light-weight, and high quality. We present a framework for the exploration of large datasets stored on a remote server with a thin client that is capable of rendering and managing full quality volume visualizations.

# Progressive Volume Rendering of Large Unstructured Grids

Steven P. Callahan<sup>1</sup>

Louis Bavoil<sup>1</sup>

Valerio Pascucci<sup>2</sup>

Cláudio T. Silva<sup>1</sup>

<sup>1</sup> Scientific Computing and Imaging Institute, University of Utah

<sup>2</sup> Lawrence Livermore National Laboratory




Figure 1: A sequence of progressive volume rendering steps for the SF1 dataset with about 14 million tetrahedra. Starting from an interactive mode that uses only the boundary (left), the algorithm progressively refines the image using incoming geometry as well as the results of the previous refinement until the full-quality rendering is achieved (right).

## ABSTRACT

We describe a new progressive technique that allows real-time rendering of extremely large tetrahedral meshes. Our approach uses a client-server architecture to incrementally stream portions of the mesh from a server to a client which refines the quality of the approximate rendering until it converges to a full quality rendering. The results of previous steps are re-used in each subsequent refinement, thus leading to an efficient rendering. Our novel approach keeps very little geometry on the client and works by refining a set of rendered images at each step. Our interactive representation of the dataset is efficient, light-weight, and high quality. We present a framework for the exploration of large datasets stored on a remote server with a thin client that is capable of rendering and managing full quality volume visualizations.

**Keywords:** Volume Rendering, Large Unstructured Grids, Client-Server, Progressive Rendering, Level-of-Detail

## 1 INTRODUCTION

Unstructured grids (*e.g.*, tetrahedral meshes) are common in the simulation domains such as computational fluid dynamics and structural mechanics. Because of the constantly improving algorithms and computational power, the scale of these simulations has continued to grow at a rate faster than visualization techniques can explore them. Specialized graphics clusters have been developed to visualize large datasets in parallel and on large displays. However, the availability of these clusters is often limited. More recently, many techniques have been developed to visualize volumetric data on commodity PCs using graphics hardware [6, 29]. This provides a solution that allows researchers to perform their visualizations locally when other resources are unavailable. However,

due to the limitations of storage and memory with most desktop machines or laptops, this solution does not scale well for extremely large datasets.

As an example, consider a scientist working remotely who would like to visualize a large dataset on his laptop computer. A reduced representation of the data (*e.g.*, simplification [13]) may not be appropriate if a high quality visualization is required for analysis. Complicating matters even further, the laptop may not have capacity on the hard disk or in memory to keep the dataset. The problem is compounded if you consider that the scientist may only want to browse through a series of datasets quickly, requiring the download of each dataset before visualization.

In this paper, we present a client-server architecture for hardware-assisted, progressive volume rendering. The main idea is to create an effect similar to progressive image transmission over the internet. A server becomes a data repository and a client (*i.e.*, a laptop with programmable graphics hardware) becomes a renderer that accumulates the incoming geometry and displays it in a progressively improving manner. Our progressive strategy is unique because it only requires the storage of a few images on the client for the incremental refinement. For interactivity, a small portion of the mesh is stored on the client which uses a bounded amount of memory. Furthermore, the progressive representation provides a natural means for level-of-detail exploration of large dataset without an explicit simplification step which may be difficult and costly for a very large dataset. Because the geometry is rendered in steps, the user can stop a progression and change the view without penalty, thus facilitating exploration. Our algorithm is robust, memory efficient, and provides the ability to create and manage approximate and full quality volume renderings of unstructured grids too large to render interactively at full resolution.

Our contributions include:

- We introduce a client-server architecture and interface for rendering large datasets and managing the resulting visualizations;
- We detail a server that acts as a data repository by streaming

a tetrahedral mesh in partial visibility order to one or more clients;

- We describe a client that uses hardware-assisted, progressive volume rendering to provide an interactive approximation, progressive refinement, and full-quality rendering of large datasets;
- We provide experimental results of our algorithm and include a discussion on the benefits and limitations of our approach.

The rest of the paper is outlined as follows. Section 2 provides an overview of related research. Section 3 describes an overview of our client-server architecture. More detail is provided about the server in Section 4 and about the client in Section 5. In Section 6, we outline our experimental results, in Section 7 we include a discussion of the trade-offs of our algorithm, and in Section 8, we provide brief discussion of our algorithm and our conclusions.

## 2 PREVIOUS WORK

**Hardware-Assisted Volume Rendering.** Volume rendering on a commodity PC has been the subject of much research recently, due to the steady increase in processing power on graphics processing units (GPUs) and the advent of programmable shaders. Here we summarize the state-of-the-art for hardware-assisted volume rendering of unstructured grids; for more detail we refer the reader to a recent survey [27]. For unstructured grids, volume rendering algorithms generally fall into two categories: ray-casting and splatting. Recently, Weiler *et al.* [29] proposed an algorithm to perform ray-casting completely on the GPU by storing the mesh and traversal structure in GPU memory. This algorithm was made more efficient and extended to handle non-convex meshes in subsequent work by Weiler *et al.* [30] and Bernardon *et al.* [2]. These algorithms benefit from low latency because they avoid CPU to GPU transfer of data. However, the limited memory of GPUs prevents these algorithms from rendering even moderately sized datasets. Pioneering work on tetrahedral splatting by Shirley and Tuchman [26] introduced the Projected Tetrahedra (PT) algorithm. For each viewpoint, PT decomposes a tetrahedron into one or more triangles that can be rendered efficiently in hardware. Unfortunately, compositing of the triangles requires an explicit visibility ordering that is implicit to ray-casting. Many algorithms have been proposed to perform the visibility ordering in object-space [32]. More recent work by Callahan *et al.* [6] introduced the Hardware-Assisted Visibility Sorting (HAVS) algorithm which sorts in image-space and object-space. The HAVS algorithm is fast, efficient, and is flexible enough to handle dynamically changing geometry [1]. In this paper, we exploit these benefits to provide a client renderer that handles streaming geometry in a progressive form. A more detailed overview of the algorithm is provided in Section 5.

**Level-of-Detail.** Hardware-accelerated volume rendering algorithms allow the interactive rendering of moderately sized datasets. However, multiresolution or level-of-detail (LOD) techniques [21] are still essential to render large datasets interactively. Unlike structured grids, the area is relatively new and not well studied for unstructured grids. For this reason, initial algorithms applied techniques for structured grids on unstructured grids by sampling them uniformly [19]. A drawback with this type of approach is that the resulting resampled meshes are larger than the original by at least an order of magnitude. A more efficient algorithm proposed by Museth *et al.* [23] uses a hierarchy of splats to render culled, opaque geometry that can be explored with the use of CSG cuts. For direct volume rendering, Cignoni *et al.* [9] present a hierarchy of simplification steps that can be traversed to a desired LOD for interactive visualization. In general, these algorithms are not

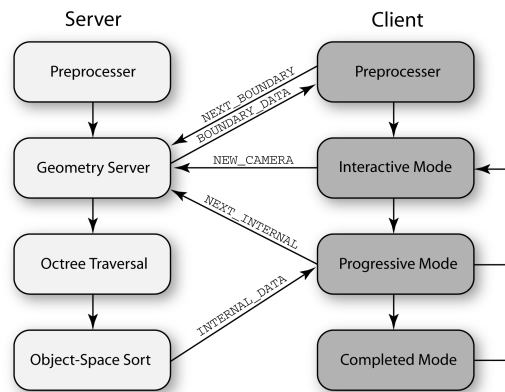


Figure 2: The client-server architecture. Communication between the client and the server is shown with annotated arrows.

suitable for streaming in a client-server architecture due to a combination of increased data size, inability to display a full-quality volume rendering, or the difficulty of traversing the levels-of-detail dynamically. More recently, Callahan *et al.* [5] introduced sample-based simplification for dynamic LOD on unstructured grids. The algorithm samples the geometry by importance and determines the number of elements to render at each pass based on the frame-rate of the previous pass. Our algorithm uses a similar idea to create an approximate volume rendering from the portion of the geometry that has arrived from the server.

**Client-Server Visualization.** The problem of remotely visualizing large datasets has been the subject of research for many years. The most widely recognized solutions perform the visualization task on large clusters using software algorithms [25] or with hardware-assisted algorithms [22] through the use of specialized graphics hardware [14]. Typically, an image is created using the cluster, then compressed for transmission to the client, where it is decompressed and displayed to the user [11]. Systems such as Vizserver<sup>1</sup> are available from vendors for performing client-server visualization in this manner. The Visapult system [3] was developed to push more of the burden of the client. This is done by rendering blocks of the data from the server in a distributed system and compositing the results on the client. A less restrictive class of algorithms performs the visualization on more limited resources by assuming a simple server (*e.g.*, a web server). To this end, Lippert *et al.* [20] introduced a system in which the server stores compressed wavelet splats that are transmitted to the client for rendering. As the splats are received by the client, the image is progressively refined. Another approach by Engel *et al.* [10] describes a progressive isosurface visualization algorithm for use on the WWW. This is done by allowing the server to compute a hierarchy of isosurfaces that are transmitted to the client progressively. For efficiency, only the difference between two successive levels of the hierarchy is sent across the network. More recently, the client-server architecture provided by Kaehler *et al.* [18] performs visualization of Adaptive Mesh Refinement (AMR) data that is stored remotely on a server and adaptively rendered locally on a client by interpolating the hierarchical structure of the grids. Our algorithm is similar to this latter class of algorithms, but for the more difficult case of unstructured grids. We assume a limited server that prepares the geometry and streams it to the client in a series of progressive steps that avoid redundant transmission and unnecessary storage. This allows the client to receive the non-overlapping geometry and refine the image with assistance of the GPU.

<sup>1</sup><http://www.sgi.com/products/software/vizserver>

### 3 SYSTEM OVERVIEW

Our client-server system architecture is depicted in Figure 2. The server acts as a data repository and geometry processor. The data is stored on the server hierarchically for efficient traversal and iterative object-space sorting. The client keeps the boundary triangles for an interactive mode and requests geometry from the server in a progressive mode. In the progressive mode, the client uses hardware-assisted LOD volume rendering to refine the image using the results of the previous progressive step. Upon completion of the progressive volume rendering, the client saves a copy of the image for later browsing in a completed mode. Viewing changes in any of the client steps causes the progressive renderer to stop and return to the interactive mode.

### 4 THE SERVER

#### 4.1 Geometry Processing

**Overview.** The server acts as a data repository that sorts and streams triangles. However, sorting large datasets in one pass may be cumbersome. For the client to remain interactive, it should begin receiving nearly-sorted faces immediately from the server. Furthermore, the client should be able to interrupt the streaming of the faces at any time to keep the exploration interactive.

Our solution is to use the server to process the faces of the mesh into an octree structure that can be traversed by depth ranges, from front to back. For every packet of faces requested from the client, the server first culls faces outside the current depth range, sorts the remaining faces, and sends them to the client. Subsequent packets use incremented depth ranges. This has the effect of distributing the sorting burden to each step of the progression. It also prevents unused geometry from being sorted.

**Depth-range culling.** For a given depth range, we cull the associated geometry using a depth-range octree. Our octree is a geometric partition of the faces, according to their centroids. Our depth-range octree is similar to the octree from [31] for isosurface extraction. However, instead of using scalar values, our octree uses dynamically changing depth ranges to cull the geometry outside the current depth range. Each octree node contains a vector of face indices. The depth range of a node is the minimum and maximum distance from the eye to the bounding box of the node. To find the faces matching a given depth range, the octree nodes are culled hierarchically by traversing nodes that may contain triangles in the given depth range. Next, the collection of triangles in the matching octree nodes are culled according to their distance from the viewpoint. The remaining triangles are then inserted into a vector for sorting. We use a radix sort on the face centroids, as in [6]. The triangles are then sent as a vertex array for direct rendering on the client.

To increment the depth range on each pass, we uniformly divide the range of the minimum and maximum distance from the eye to the bounding box of the mesh. This has the unfortunate side effect that the number of triangles per slice can vary in size. Since the sorting and the network transfers are more efficient for large packets, we collect packets and only transmit them to the client when a user-specified target is reached.

**View-frustum culling.** When interactively exploring regions of the mesh in detail, often many of the faces are outside of the view-frustum. These faces should not be transmitted to the client. Therefore, in addition to the depth range test, we also perform view-frustum culling for each node of the octree. The left, right, bottom and top planes of the frustum are computed from the modelview-projection matrix sent by the client. This results in significant performance improvements for zoomed-in views of the dataset.

**Compression.** Network transfers become a bottleneck for client-server systems with high bandwidth. To reduce the bandwidth, we compress the transmitted vertex arrays using the open-source zlib library<sup>2</sup> based on Huffman codes, which is fast and robust. We use the maximum level of compression in all our results.

#### 4.2 Network Protocol

The following description assumes a single client per server, and can be extended for multiple clients by storing the context of the stream on each client.

The server understands three types of commands: `NEW_CAMERA`, `NEXT_BOUNDARY`, and `NEXT_INTERIOR`.

**New Camera.** Each new client initializes the stream by sending an `NEW_CAMERA` command, containing a frame ID and the camera information. The frame ID is an integer initialized to 0 and incremented by the client for every `NEW_CAMERA` command.

**Next Boundary.** Upon connection, the client requests boundary faces from the server using the `NEXT_BOUNDARY` command. The server incrementally sends the boundary faces in a `BOUNDARY_DATA` packet consisting of an unsorted triangle soup (*i.e.*, a sequence of vertex coordinates) for the client to use during interactive rendering. By sending the boundary in chunks, the client can request only the portion that it can retain in memory. If the client is limited and can only use a portion of the boundary during interaction, the `NEXT_BOUNDARY` command can be used at the beginning of the progressive rendering to fill in the missing boundaries before the internal faces are transmitted.

**Next Interior.** To avoid unnecessary sorting when the camera is moving, the server does not sort the faces until it receives a `NEXT_INTERIOR` command from the client. The server culls the geometry by depth and frustum, sorts by centroid, and sends back a chunk of the interior faces in an `INTERIOR_DATA` packet.

**Synchronization issues.** We use a TCP/IP socket to transmit our data. Unlike UDP, TCP guarantees that the data packets sent from the server arrive in the same order they were sent and without error.

Since the client can change the camera at any time, the state of the client and the server need to be synchronized. This is necessary to avoid issues when the camera moves and the client is receiving data asynchronously, or the server is processing data. All the packets from the client or from the server contain a frame ID in their header. Before processing a packet, the client and the server check that the frame ID from the packet is the same as the current frame ID. Packets with obsolete frame IDs are ignored by the server and the client. In addition, when the client receives an obsolete packet, it resets the stream.

### 5 THE CLIENT

The main function of the client is as a progressive volume renderer. Because the client may be limited in disk space as well as memory (*e.g.*, a laptop), the goal is to avoid keeping information on the machine at each progressive step. Therefore, the client acts as a *stream renderer*— it receives geometry transmitted from the server and renders it directly with the GPU. This requires a volume renderer capable of handling dynamic data in an efficient manner. In addition, our algorithm requires an approximation technique for partial geometry as well as a means of keeping previously computed information for subsequent refinement steps.

To leverage the GPU efficiently, we extend the Hardware-Assisted Visibility Sorting (HAVS) algorithm of Callahan *et al.* [6] to perform progressive volume rendering. The HAVS algorithm

<sup>2</sup><http://www.zlib.net>

operates in both object-space and image-space to sort and composite the triangles that compose a tetrahedral mesh. In object-space, the triangles are sorted by their centroids using a linear-time radix sort for floating-point numbers. This provides a partial order for the triangles. Upon rasterization, the fragments are sorted again in image space using a fixed size A-buffer [7] implemented with programmable shaders called the  $k$ -buffer. For each pixel of the resulting image,  $k$  entries (scalar value  $v$  and depth  $d$ ) are stored in textures on the GPU. An incoming fragment is compared to entries in the  $k$ -buffer to find the two closest to the current viewpoint (for front-to-back compositing). These entries are then used to look up the color and opacity for the volume gap in a pre-integrated table by using the front scalar, back scalar, and distance between the entries. This color and opacity are then composited to the framebuffer, the front entry is discarded, and the remaining entries are written back into the  $k$ -buffer. For more detail on the HAVS algorithm, see [6]. Our progressive renderer uses the HAVS algorithm as a basis for sorting and compositing. The server handles the object-space sorting, while the client handles the fixed-size image-space sorting and compositing.

The  $k$ -buffer is implemented in hardware using multiple render targets (MRTs), which allows the reading and writing of multiple off-screen textures in each pass. Currently, hardware limits the number of MRTs to four, which limits the size of  $k$  to six (one texture for an off-screen framebuffer and three textures for  $k$ -buffer entries). In OpenGL, the simplest access to multiple render targets is in the form of Framebuffer Objects (FBOs). An FBO is a collection of logical buffers such as color, depth, or stencil. Multiple color buffers (up to four) can be *attached* to an FBO for off-screen rendering. FBOs make it possible (and efficient) to swap attached buffers between rendering passes. Currently, this is faster than switching between multiple FBOs. The ability to render into a subset of buffers in multiple passes is at the heart of our progressive algorithm.

The progressive volume rendering works by using five different render targets for each pass, represented as four channel, 32-bit floating-point textures. The textures are used as follows:

$T_{pro}$ : An off-screen framebuffer for the progressive image ( $R_p, G_p, B_p, A_p$ ).

$T_{k12}$ :  $k$ -buffer entries 1 and 2 ( $v_1, d_1, v_2, d_2$ ).

$T_{k34}$ :  $k$ -buffer entries 3 and 4 ( $v_3, d_3, v_4, d_4$ ).

$T_{k56}$ :  $k$ -buffer entries 5 and 6 ( $v_5, d_5, v_6, d_6$ ).

$T_{approx}$ : A temporary framebuffer for the approximation of the portion of the mesh not yet received ( $R_a, G_a, B_a, A_a$ ).

A combination of these textures is used for each step of the progressive volume rendering. The contents of all the textures except  $T_{approx}$  are reused in subsequent progressive passes.

Our progressive volume rendering is separated into three modes of operation. *Interactive Mode* is used during camera events such as rotation, pan, or zoom. *Progressive Mode* is used when interaction stops to stream triangles from the server to the client in chunks. The *Progressive Mode* can be interrupted at any time if the user begins interaction again or the stream finishes. When a complete image is generated with the *Progressive Mode*, *Completed Mode* automatically stores the image for future browsing.

### 5.1 Interactive Mode

A reduced representation of the original mesh is often necessary when considering large datasets. Our interactive mode has two requirements. First, it is fast enough to render at interactive rates (e.g., 10 fps). Second, the results of the interactive mode can be




Figure 3: The progressive volume render on the client. A progressive buffer is maintained between steps and an approximate buffer is created to fill the unknown region of the mesh.

used as a first step in the progressive volume rendering. Because of the second requirement, level-of-detail techniques that require simplification hierarchies [9] are not feasible. Instead, we use a similar approach to Callahan *et al.* [5], where a subset of the original geometry is used to create a reduced representation. In particular, Callahan *et al.* noticed that an efficient approximation of the dataset can be created by computing the volume rendering integral between only the boundary fragments of the mesh. We use a more robust version of this approximation that can be used in future progressive steps.

Upon connection with the server, the client receives some initial data to begin the progressive rendering process. First, mesh parameters such as maximum edge length and scalar range are transferred for creating a pre-integrated lookup table. Immediately following these parameters, the boundary triangles of the mesh are transferred to the client. These boundary triangles are placed in arrays on GPU memory and remain there for the duration of the client-server connection. If the entire boundary cannot be maintained in memory, a subset of the boundary can be used instead.

During user interaction (i.e., rotation, panning, or zooming), the following steps take place to create an approximate image.

1. Buffer  $T_{k12}$  is attached to the FBO as well as a depth buffer. The depth buffer is cleared and set to GL\_LESS for a first pass at the boundary geometry. The depth buffer is then cleared again and set to GL\_GREATER for a second pass at the boundary geometry. This has the same effect as depth-peeling [12] the front and back fragments and placing them in the  $k$ -buffer.
2. Buffers  $T_{approx}$  and  $T_{k12}$  are attached to the FBO and a screen-aligned plane is rendered. Color and opacity for the ray-gaps between the front and back fragments from the boundary are looked up and composited into  $T_{approx}$ .
3. Buffer  $T_{approx}$  is displayed to the screen.

These steps are repeated for every view change. When the user stops interacting, the results of  $T_{k12}$  is used for the progressive steps. The importance of using this depth-peeling approach instead of the  $k$ -buffer directly is discussed in more detail in Section 5.4.



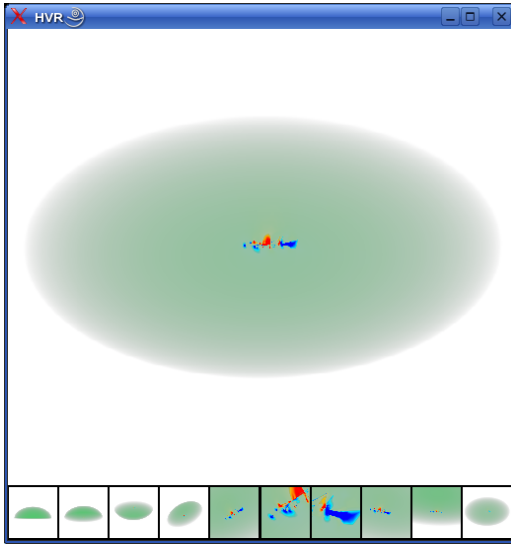


Figure 4: A snapshot of the interaction with the *Completed Mode* for the F16 dataset. Upon completion of a full quality rendering, the image is automatically stored for future browsing by selecting the icons at the bottom of the window. The user may also save intermediate steps with a keystroke.

## 5.2 Progressive Mode

After the boundary has been drawn in *Interactive Mode*, the internal triangles of the mesh are streamed in an approximate front-to-back order based on their centroids. For each portion of the geometry that arrives from the server, the progressive volume renderer takes the following steps:

1. Buffer  $T_{approx}$  is cleared.
2. Buffers  $T_{pro}$ ,  $T_{k12}$ ,  $T_{k34}$ , and  $T_{k56}$  are attached to the FBO and the incoming internal geometry is rendered. The  $k$ -buffer is used to sort the fragments and composite the results into  $T_{pro}$ . This step is similar to HAVS, but with only a portion of the geometry.
3. Buffer  $T_{approx}$  is attached to the FBO,  $T_{k12}$ ,  $T_{k34}$  and  $T_{k56}$  are bound as a read-only textures, and a screen-aligned plane is drawn. The fragment shader finds the  $k$ -buffer entry  $f$  closest to the current view and the entry  $b$  farthest from the current view, looks up the color and opacity for the ray gap between  $f$  and  $b$ , and composites the result into  $T_{approx}$ .
4. Buffers  $T_{pro}$  and  $T_{approx}$  are attached to the FBO and a screen-aligned plane is drawn. The fragment shader composites  $T_{pro}$  into  $T_{approx}$ .
5. Buffer  $T_{approx}$  is displayed to the screen.

The  $k$ -buffer entries and progressive buffer  $T_{pro}$  are then ready to use for the next rendering pass. Figure 3 illustrates a progressive step. This process is repeated until all the geometry has been streamed from the server. When rendering the last chunk of geometry from the server, an additional step is taken to flush the  $k$ -buffer entries into the  $T_{pro}$  by rendering  $k - 1$  screen-aligned planes after the second step. After which, the  $T_{pro}$  buffer contains the full quality volume rendering.




Figure 5: A zoomed-in view of the SF1 dataset. Significant performance improvements are made by frustum-culling the geometry on the server. Here, 70% of the geometry is culled.

## 5.3 Completed Mode

Once the stream of geometry has terminated and the progressive volume rendering is completed, we capture the  $T_{pro}$  buffer and save it for subsequent browsing as shown in Figure 4. Our progressive volume renderer allows the user to browse any previously completed visualizations by selecting the corresponding thumbnail. Capturing this data for the user has important benefits. First, it prevents the user from losing important visualizations through interactions that could reset the results of the previous stream. Second, it allows a user to specify a set of camera positions to the progressive volume renderer so the user can easily capture an animation of the exploration process. This tool is useful for exploring previously generated results while the current view is being progressively rendered off-screen.

## 5.4 Considerations

The  $k$ -buffer algorithm efficiently handles streaming geometry because it simultaneously reads and writes from textures at each pass. Due to the high parallel nature of GPU architectures, this may result in a race condition for overlapping triangles. Because the HAVS algorithm sends geometry sorted by centroid, it effectively layers the geometry in the depth direction and thus avoids these errors. However, when rendering only boundary geometry, this race condition can be noticeable independent of sorted order because the depth complexity is small. Because our progressive rendering quality depends on the first pass in *Interactive Mode*, these errors are not as acceptable as they are with an LOD approach [5]. Therefore, we perform the depth-peeling of the front and back fragments before inserting them into the  $k$ -buffer. This resolves the race condition, improves the quality of the rendering, and maintains interactive rates (see Section 6).

The depth-peeling as described has the unfortunate side-effect that it removes any non-convexities in the boundary. This can propagate through the progressive steps and cause the empty space to be composited into the final image. Since storing the back fragments in the  $k$ -buffer effectively reduces  $k$  by one during the progressive steps, storing all boundary information for non-convex objects can severely impact sorting capabilities. A solution to this is to transmit boundary and internal faces during progressive steps. This would introduce redundant fragments only in the front and back and al-

| Dataset   | Tetrahedra | Preprocess | Server   |       | Data Transfer | Client      |             | Total   |
|-----------|------------|------------|----------|-------|---------------|-------------|-------------|---------|
|           |            |            | Traverse | Sort  |               | Interactive | Progressive |         |
| Fighter   | 1.4 M      | 20 s       | 2.0 s    | 0.1 s | 23.3 s        | 0.01 s      | 0.5 s       | 25.9 s  |
| F16       | 6.3 M      | 118 s      | 3.6 s    | 0.9 s | 144.8 s       | 0.01 s      | 2.7 s       | 152.0 s |
| F16 (25%) | 6.3 M      | 118 s      | 1.16 s   | 0.1 s | 22.6 s        | 0.01 s      | 2.3 s       | 26.2 s  |
| SF1       | 13.9 M     | 447 s      | 47.4 s   | 1.4 s | 529.5 s       | 0.01 s      | 4.0 s       | 582.3 s |
| SF1 (25%) | 13.9 M     | 447 s      | 6.8 s    | 0.3 s | 53.7 s        | 0.01 s      | 3.5 s       | 64.3 s  |

Table 1: Performance analysis of the preprocessing and one step of the progressive volume rendering over a network.

low other boundary fragments to be used in the progression steps to avoid compositing empty space, as in [6]. To completely remove the storage overhead of the back boundary in the  $k$ -buffer, an extra texture can be used to store the back fragments during *Interactive Mode* and can be bound as a read-only texture during *Progressive Mode*.

## 6 RESULTS

Our experimental results were measured on a client running Linux with two Dual core Opteron 2.25 GHz processors, an NVidia GeForce 7800 GTX graphics card, and 4 GB RAM. The server machine was also running Linux with 3.0 GHz Pentium D processor and 3 GB RAM. Performance timings are measured with a  $512 \times 512$  viewport on a 100 Megabit/sec network with regular network load. Our experiments include timing results for the rendering and client-server architecture, as well as error measurements for the progressive images.

Our timing measurements are shown in Table 1 for three meshes of varying size. The measurements can be broken into four important sections: preprocessing, the server, the client, and data transfer from the client to the server. The preprocessing step occurs on the server and includes loading the file from disk, separating boundary from internal triangles, building an octree, and transferring the initial data to the client. It is important to note that the octree can be built once and stored for efficient loading, thus avoiding most of the preprocessing. The timing results for the server, client, data transfer, and total time represent one progressive rendering of the dataset from views that include the whole mesh as well as views showing a zoomed-in portion (showing 25%) which benefits from frustum culling (see Figures 5, 7, and 8). In our experiments, the interactive portion of our progressive renderer achieved at least 100 frames-per-second for every dataset, thus will scale well for even larger datasets. The transferring of the data from the client to the server is the bottleneck of the system and takes about 90% of the total time for one visualization. Though the focus of this paper is a client-server architecture for progressive rendering, our algorithm also can be used as a data exploration tool running locally on a workstation. When running locally, you lose the advantage of a data repository, but effectively remove the data transfer time, which improves the interactivity substantially.

| Stream Size    | Octree Depth 7 | Octree Depth 8 |
|----------------|----------------|----------------|
| 1K Triangles   | 0.096 s        | 0.119 s        |
| 10K Triangles  | 0.194 s        | 0.339 s        |
| 100K Triangles | 0.919 s        | 1.092 s        |

Table 2: Latency analysis of different server settings on the Fighter dataset.

An important consideration for a client-server architecture is the latency of the network. Several settings on the server effect the round-trip latency of the system — the time for the client to send a packet and receive a response. An obvious consideration is the compression of the geometry during transmission. Our experiments

show an average of about a 50% compression rate for all parameters tested. Other important considerations include the octree resolution and the stream size or the number of triangles sent on each progressive step. Table 2 shows the effect of these parameters on the latency for the Fighter dataset. In our performance analysis of Table 1 we use 10K triangles per octree node and a stream size of 10K triangles.

Our final experiment was to analyze the quality of the progressive steps. To measure this metric, we used root mean squared (RMS) error to compare incremental steps with the final rendering for all of our experimental datasets. Figure 6 shows a plot of these errors as the progression refines. These results show that the image quality steadily converges to the full quality image, which is important for allowing the user to explore the dataset efficiently.

## 7 DISCUSSION

Our progressive volume rendering system is unique because it efficiently handles data of arbitrary size on a thin client. This is done using a novel technique which keeps only image data and boundary geometry in GPU memory on the client for each step. The amount of memory used on the client can be bound by adjusting the stream size and the size of the boundary. In our experiments, the boundary was not large enough to adversely effect performance or expend memory constraints. In fact, even with a 14 million tetrahedra dataset, the boundary can be volume rendered with our algorithm in *Interactive Mode* at over 100 frames-per-second. Even if a dataset has a boundary too large to fit in GPU memory, our algorithm would work efficiently by using only a random subset of the boundaries for an approximate rendering in interactive mode. This would have the effect of lightening the general appearance of the approximation. The remaining boundary triangles could then be streamed before the rest of the internal faces.

A client-server progressive volume renderer is advantageous because it allows the data to be stored in a central repository, while the rendering can be performed with the help of graphics hardware on a client. This efficiently splits the load of the client and server for datasets that may be too large to render locally. Since processing power continues to increase rapidly for both CPUs and GPUs, it is becoming increasingly important to develop algorithms that efficiently harness the computational power without being limited by memory constraints. Our algorithm uses this approach by acting as a stream renderer and allowing the interactive exploration of a dataset with only a portion of the geometry.

The main disadvantage of using this paradigm is that data transfer over the network incurs a substantial penalty. We reduce this penalty partially with the use of lossless compression of the stream. If some loss in quality is acceptable, quantization techniques could also be applied to reduce the bandwidth of the geometry. Because data transfer is a limiting factor in quickly visualizing a full quality rendering, the quality of the approximation is important. Unlike naive approaches that render only an opaque boundary mesh or outline, our initial approximation gives excellent results with little overhead. With only a few iterations, our progressive volume renderer converges to the final image which facilitates dynamic explo-




Figure 6: Root Mean Squared Error (RMS) of the progressive images going from only the boundaries (0%) to the full quality (100%).

ration. We find this aspect useful because often in the exploration process, the user will not wait for the entire progressive volume rendering before moving on to another viewpoint. Because only the geometry within the current view frustum is transferred, efficiently exploring details of the dataset becomes easy. When combined with the image capturing system, these features provide a powerful data exploration tool for large datasets.

## 8 CONCLUSIONS AND FUTURE WORK

Our algorithm provides a progressive volume rendering system for interactively exploring large datasets. We use a novel approach of balancing the load of the client and server while minimizing the memory constraints of the client. In fact, our algorithm can bound the memory usage of the client machine to allow a wide variety of client devices. We have introduced a novel progressive algorithm that efficiently uses the GPU to incrementally refine the visualization by retaining only image data. An interactive mode can be efficiently computed with the addition of boundary triangles that can be kept in GPU memory. To further improve interaction, our system keeps previously computed visualizations that can be interactively browsed while progressive rendering the visualization. Finally, we have provided detailed experiments and discussed the trade-offs of a client-server approach for volume rendering.

In the future, we would like to explore more efficient methods of reducing the data size for transmission from server to client. In particular, geometry compression algorithms (*i.e.*, [4]) could alleviate bandwidth constraints. Another area of future research is to extend our system to render isosurfaces and to explore the visualization of time-varying data in a progressive manner.

## ACKNOWLEDGMENTS

The authors acknowledge Neely and Batina (NASA) for the Fighter dataset, Tremel (EADS-M) for the F16 dataset, and O'Hallaron and Shewchuk (CMU) for the SF1 dataset. Steven P. Callahan and Louis Bavoil are supported by the Department of Energy (DOE) under the VIEWS program. Cláudio T. Silva is partially supported by the DOE under the VIEWS program and the MICS office, the National Science Foundation (grants CCF-0401498, EIA-0323604, OISE-0405402, and IIS-0513692), and a University of Utah Seed Grant.




Figure 7: A zoomed-in portion of the F16 dataset (about six million tetrahedra) shown progressively at (a) 0%, (b) 12%, (c) 61%, and (d) 100%.

## REFERENCES

- [1] F. F. Bernardon, S. P. Callahan, J. L. D. Comba, and C. T. Silva. Volume rendering of unstructured grids with time-varying scalar fields. In *EG Symposium on Parallel Graphics and Visualization*, 2006. to appear.
- [2] F. F. Bernardon, C. A. Pagot, J. L. D. Comba, and C. T. Silva. GPU-based tiled ray casting using depth peeling. *Journal of Graphics Tools*, to appear. Also available as SCI Institute Technical Report UUSCI-2004-006.
- [3] W. Bethel, B. Tierney, J. Lee, D. Gunter, and S. Lau. Using high-speed wans and network data caches to enable remote and distributed visualization. In *Supercomputing '00*, page 28, 2000.
- [4] U. Bischoff and J. Rossignac. Tetstreamer: Compressed back-to-front transmission of delauney tetrahedra meshes. In *Data Compression Conference*, 2005.
- [5] S. P. Callahan, J. L. D. Comba, P. Shirley, and C. T. Silva. Interactive rendering of large unstructured grids using dynamic level-of-detail. In *IEEE Visualization '05*, pages 199–206, 2005.
- [6] S. P. Callahan, M. Ikits, J. L. D. Comba, and C. T. Silva. Hardware-assisted visibility ordering for unstructured volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 11(3):285–295, 2005.
- [7] Loren Carpenter. The A-buffer, an antialiased hidden surface method. In *Computer Graphics (Proceedings of ACM SIGGRAPH)*, volume 18, pages 103–108, July 1984.
- [8] P. Chopra and J. Meyer. Incremental slicing revisited: Accelerated volume rendering of unstructured meshes. In *IASTED Visualization, Imaging, and Image Processing*, pages 533–538, 2002.
- [9] P. Cignoni, L. De Floriani, P. Magillo, E. Puppo, and R. Scopigno. Selective refinement queries for volume visualization of unstructured tetrahedral meshes. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):29–45, 2004.
- [10] K. Engel, O. Sommer, C. Ernst, and T. Ertl. Progressive isosurfaces on the web. Late Breaking Hot Topics, IEEE Visualization '98, 1998.
- [11] K. Engel, O. Sommer, and T. Ertl. A framework for interactive hardware accelerated remote 3D-visualization. In *EG/IEEE TCVG Sym-*






Figure 8: A zoomed-in portion of the Fighter dataset (about 1.4 million tetrahedra) shown progressively at (a) 0%, (b) 14%, (c) 29%, (d) 43%, (e) 57%, (f) 71%, (g) 86%, and (h) 100%. View-frustum culling on the server removed 75% of the original geometry for this view.

- posium on Visualization (VisSym), 2000.
- [12] C. Everitt. Interactive order-independent transparency. White paper, NVIDIA Corporation, 1999.
  - [13] M. Garland and Y. Zhou. Quadric-based simplification in any dimension. *ACM Transactions on Graphics*, 24(2), April 2005.
  - [14] G. Humphreys, M. Houston, R. Ng, R. Frank, S. Ahern, P. D. Kirchner, and J. T. Klosowski. Chromium: a stream-processing framework for interactive rendering on clusters. In *SIGGRAPH '02*, pages 693–702, 2002.
  - [15] M. Isenburg and P. Lindstrom. Streaming meshes. In *IEEE Visualization 2005*, pages 231–238, oct 2005.
  - [16] M. Isenburg, P. Lindstrom, S. Gumhold, and J. Shewchuk. Streaming compression of tetrahedral volume meshes. In *Graphics Interface*, 2006. to appear.
  - [17] M. Isenburg, P. Lindstrom, S. Gumhold, and J. Snoeyink. Large mesh simplification using processing sequences. In *Visualization'03 Conference Proceedings*, pages 465–472, 2003.
  - [18] R. Kaehler, S. Prohaska, A. Hutanu, and H.-C. Hege. Visualization of time-dependent remote adaptive mesh refinement data. In *IEEE Visualization '05*, pages 175–182, 2005.
  - [19] J. Leven, J. Corso, J. D. Cohen, and S. Kumar. Interactive visualization of unstructured grids using hierarchical 3d textures. In *Proceedings of IEEE Symposium on Volume Visualization and Graphics*, pages 37–44, 2002.
  - [20] L. Lippert, M. H. Gross, and C. Kurmann. Compression domain volume rendering for distributed environments. *Computer Graphics Forum*, 16(3):C95–C107, 1997.
  - [21] D. Luebke, M. Reddy, J. Cohen, A. Varshney, B. Watson, and R. Huebner. *Level of Detail for 3D Graphics*. Morgan-Kaufmann Publishers, 2002.
  - [22] K. Moreland, B. Wylie, and C. Pavlakos. Sort-last parallel rendering for viewing extremely large data sets on tile displays. In *Proceedings of the IEEE 2001 symposium on parallel and large-data visualization and graphics*, pages 85–92, 2001.
  - [23] K. Museth and S. Lombeyda. Tetsplat: Real-time rendering and volume clipping of large unstructured tetrahedral meshes. In *Proceedings of IEEE Visualization 2004*, pages 433–440, 2004.
  - [24] A. Norton and A. Rockwood. Enabling view-dependent progressive volume visualization on the grid. *IEEE Computer Graphics and Applications*, 23(2):22–31, 2003.
  - [25] S. Parker, P. Shirley, Y. Livnat, C. Hansen, and P.-P. Sloan. Interactive ray tracing for isosurface rendering. In *Proceedings of IEEE Visualization '98*, pages 233–238, 1998.
  - [26] P. Shirley and A. Tuchman. A polygonal approximation to direct scalar volume rendering. *Proceedings of San Diego Workshop on Volume Visualization*, 24(5):63–70, 1990.
  - [27] C. T. Silva, J. L. D. Comba, S. P. Callahan, and F. F. Bernardon. A survey of GPU-based volume rendering of unstructured grids. *Brazilian Journal of Theoretic and Applied Computing (RITA)*, 12(2):9–29, 2005.
  - [28] H. T. Vo, S. P. Callahan, N. Smith, C. T. Silva, W. Martin, D. Owen, and D. Weinstein. iRun: Interactive rendering of large unstructured grids. Technical report, Scientific Computing and Imaging Institute, 2006.
  - [29] M. Weiler, M. Kraus, M. Merz, and Thomas Ertl. Hardware-based ray casting for tetrahedral meshes. In *IEEE Visualization '03*, pages 333–340, October 2003.
  - [30] M. Weiler, P. N. Mallón, M. Kraus, and T. Ertl. Texture-Encoded Tetrahedral Strips. In *Proceedings of Symposium on Volume Visualization 2004*, pages 71–78. IEEE, 2004.
  - [31] J. Wilhelms and A. Van Gelder. Octrees for faster isosurface generation. *ACM Transaction on Graphics*, 11(3):201–227, 1992.
  - [32] P. L. Williams. Visibility-ordering meshed polyhedra. *ACM Transactions on Graphics*, 11(2):103–126, April 1992.
  - [33] J. Wu and L. Kobbelt. A stream algorithm for the decimation of massive meshes. In *Graphics Interface '03*, pages 185–192, 2003.