Constant Rate Visualization of Large Anatomical Models for Medical Guidance

B. Takacs, G. Szijarto, D. Hanak

Virtual Human Interface Group, MTA SZTAKI, Budapest, Hungary.

Abstract

This paper describes a volumetric visualization technique designed to deliver constant rendering speed of highly detailed anatomical structures at interactive rates using portable, and low-cost computers. Our solution represents the torso section of the human body as a volumetric data set and employs label maps as the prime data format of storing anatomical structures. Multi-channel 3D textures are uploaded to the GPU and a simple pixel shader algorithm allows operators to select structures of interest in real-time. The visualization module described herein was successfully integrated into a 3D Anatomical Guidance System for Ultrasound Operators and its rendering performance tested on a "backpack" system.

Image guided surgery, ultrasound 3D guidance, human modelling, Virtual Human Interface (VHI)

---

1. Introduction

Anatomical Guidance (AG) combines generic 3D models of the human body with medical real-time imaging devices, such as laparascopic or hand-held ultrasound [1]. The purpose of image-guided interventions is to visually aid medics and operators in correctly identifying anatomical structures for the purposes of examination and planning of subsequent surgical interventions. It has been shown, that such tools doubled the performance of novice users, during the in vivo examination of anesthetized pigs, effectively boosting their skills to reach the accuracy of experts without the AG support. Furthermore experts also performed nearly 150% better when using the AG system in comparison to performing diagnostics without it [2]. These studies indicate that a portable, anatomically guided ultrasound system would find many applications in the field of patient care.

One of the key challenges to address in the design of such systems, however, lies in the difficulty of developing visualization algorithms that run on low-cost computer platforms without compromising the necessary details of the anatomical structures presented to the user. Classical segmentation and display methodology requires the construction of triangularized surfaces, which in turn are rendered by hardware accelerated polygon pipeline in the GPU. The problem with such an approach, however, lies in its rendering performance or speed being dependent on the polygon count, i.e. the level of detail of these models. The more detailed these models are or the more of the structures one needs to visualize the slower the rendering frame rate (fps) becomes.

To overcome this difficulty we present herein a rendering solution that displays segmented anatomical structures of arbitrary complexity at constant frame rates and real-time performance. To achieve this improvement we rely on the pixel-shader pipeline of the GPU and treat data structures of the human body as volumetric data sets, where the information regarding individual sub-structures is stored in *label maps*.

In the remainder of this paper is organized as follows: In Section 2 we briefly introduce our guidance system in which 3D virtual human models are being used. In section 3 we describe how over 470 different high-resolution anatomical structures were created and optimised for our purposes. Section 4 describes our label-map based colour volume visualization technique and demonstrates its use for displaying anatomical structures. Finally, Section 5 describes our evaluation tests and we conclude our discussion in Section 6.

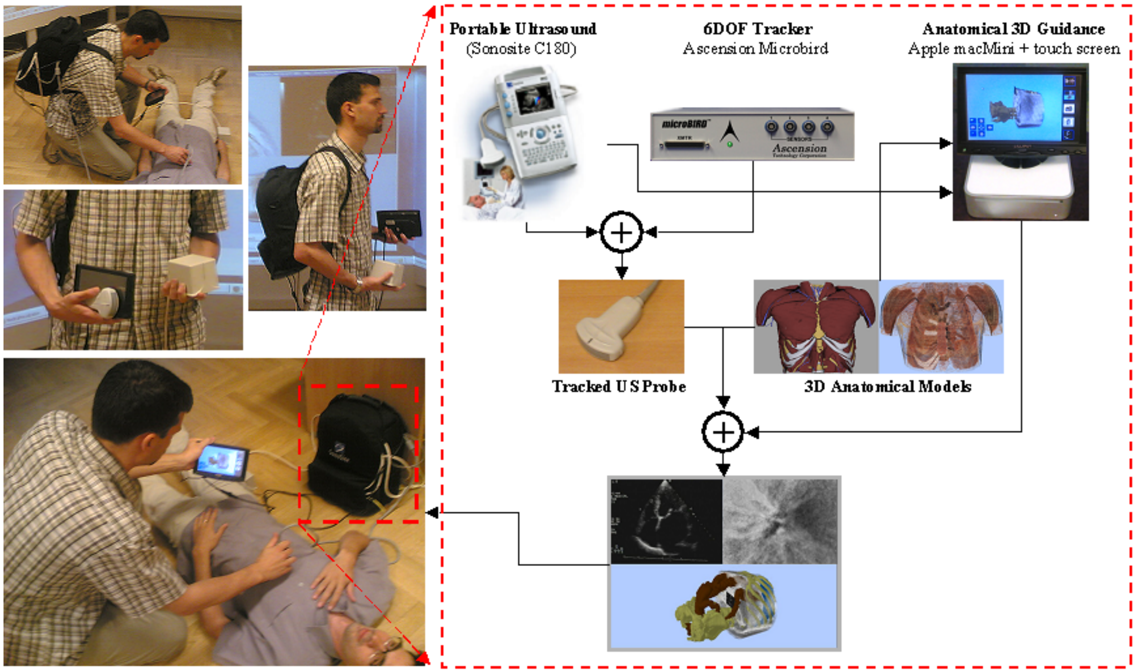

Figure 1: *A "Backpack" anatomically guided ultrasound system (Left – system in use / Right – main components).*

2. Portable Anatomically Guided Ultrasound System

The basic setup of the compact anatomically guided Ultrasound system is shown in Figure 1. On the left, a medic is using the system to detect internal bleeding of a patient lying on the ground. The entire system is carried in a backpack on his back, his hands hold three outside components, namely the US probe, the reference cube for the tracking system, and the touch-sensitive computer screen which serves as the main interface to access the system's functions. On the right the HW elements packed tightly inside the backpack are shown. From left to right a Portable Ultrasound device (Sonosite C180) is combined with a 6DOF, high precision tracker (Ascension Microbird) to create a tracked sequence of US images. This information is transmitted to a small computer (Apple macMini), the heart of our Anatomical 3D Guidance solution, that combines this information with generic models of the human body and outputs guidance information to the medic's hand-held touch screen.

3. Creating Anatomical Models

One of the key technical elements of our backpack Ultrasound guidance solution is a highly detailed digital and animatable representation of *the human body and its anatomy*. This representation comprises a generic 3D virtual human model as well as models of internal organs segmented from volumetric data representations. The purpose of this generic 3D model is to provide guidance to the operator showing major anatomical areas of interest when patient-specific data is not available. Guidance herein refers to showing regions inside the body in order to help medics properly move the US probe and thus obtain high quality US images. Guidance information is also very important to the doctors reviewing the recorded US scans, as it provides 3-dimensional context information without which interpreting the US images would practically be impossible.

The visualization system was optimized for portable performance and it offers different methodologies for showing regions of interest to the operator. They all rely on the notion of *creating Anatomical Atlases* and use them later as reference for segmenting patient-specific CT's and visualizing target information. Instead of representing structures as polygonal objects, however, our atlas contains colored *image slices* combined with detailed *label maps*. Specifically, we focused our efforts on creating a complete model of the inside the torso section of the body and optimize it for real-time viewing and manipulation on a portable computational platform.

As the first step in this process the torso segment from the volumetric data scans of the *Visible Human Male Data* set [3] was processed with an open-source segmentation and labeling tool, called *3D Slicer* [4]. *Slicer* provides a set of automated techniques to specify regions of interest in volumetric data sets and construct 3D polygonal surfaces. These organ models typically contain up to one million polygons each, thereby making them difficult to use in a real-time system. To address this problem once these structures were exported we further optimized them with the help of professional modeling tool (Maya). Figure 2 shows examples of the named geometries including muscles, bone structure of the chest and spine, vascular structure of veins and arteries, and internal organs, etc. The entire model contains 478 identified anatomical structures in total.

Besides using 3D geometry to represent parts of the human body as described in the previous section, patient specific data (e.g. a CT or MRI) can be displayed as *volumetric* data sets directly. Specifically, we implemented a volumetric visualization node that readily accept these kinds of data sets, such as MRI and CT slices, and displays them in an integrated fashion with other elements of the virtual anatomical guidance systems. This, so called *volume node* uses 3D textures (*u*,*v*,*w*) as a basic data storage and takes advantage of the latest advances in hardware-supported rendering methodologies. However, our solution does not require the use of expensive computational platforms but rather a portable personal computer or laptop.

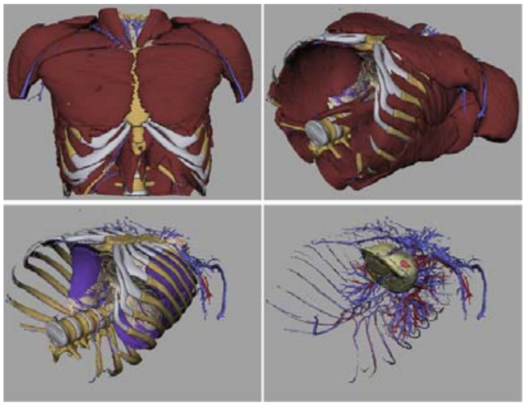

Figure 2: *Segmented and labeled anatomical structures derived from the NIH Visible Male Data set and optimized for real-time viewing and guidance.*

The volume node supports multiple ways of displaying volumetric information. As an example the volume node accepts slicing planes and can show the position of the US probe to the medic in many different ways. For flexible visualization properties, the image plane representing the ultrasound image can be used to glide through and penetrate this volume or to slice it directly Figure 3 demonstrates this concept. The figure shows the final operator display of the integrated anatomical guidance system. The upper left corner displays the live video output of a *Sonosite C180* potable ultrasound device. The upper right hand corner shows the re-sampled and processed CT volume of the patient in the section matching the 3D position and orientation of the US image plane. This information provides anatomical context even when 3D models of intestines are not readily available. The large format display below shows 3D segmented geometry (skeletal structure, colon and lung) along with the aligned CT and marker for the US plane (see center). As the medic scans the patient the position and orientation of the US slicer is updated in real-time accordingly and the volume is re-sampled and sliced as well. The advantage of the upper right hand display is that volumetric image data sets provide more readily recognizable features to the trained eye. On the other hand, for non-expert users the 3D models below offer an easy-to-understand interface for reaching their goal, i.e. to scan a specific region of the body and transmit those scans to a remote location.

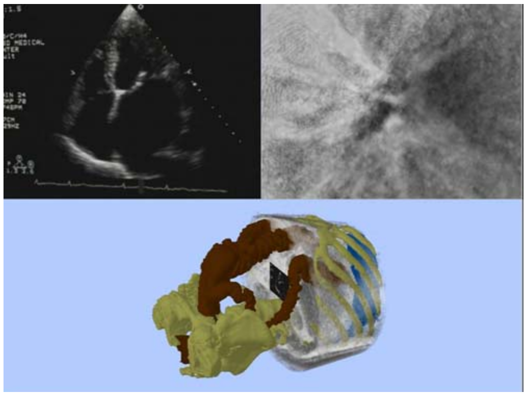

Figure 3: *Final operator display of an anatomical guidance system using segmented 3D models and volumetric textures (see text).*

While most volume architectures focus on visualizing gray scale imagery (since CT's and MRI provide density information only), we extended our visualization node to handle color volumes as well. This lays the foundation for our constant speed display algorithm described in the next section.

4. Using Label Maps for High-Speed Display

As anatomical guidance requires the visualization system to be able to show internal structures of the human body in an easily reconfigurable manner, our goal was to devise an architecture that can implement this functionality without explicit need for segmenting those structures into geometry. To address this need we developed a novel methodology that benefits from real-time image processing at the stack level as well as the parallel computing power of hardware supported pixel- and vertex shaders in the GPU. This architecture, which is called *dynamic shading*, offers a key advantage over previously existing solutions as it eliminates the need for pre-processing the data and storing large and complex 3D geometry. Instead, the pipeline - shown in Figure 4 - performs computations directly on the graphics card while the scene is being rendered, thus freeing up resources of the CPU for other calculations. The basic block diagram of dynamic shading and image processing architecture implementing this functionality is shown below.

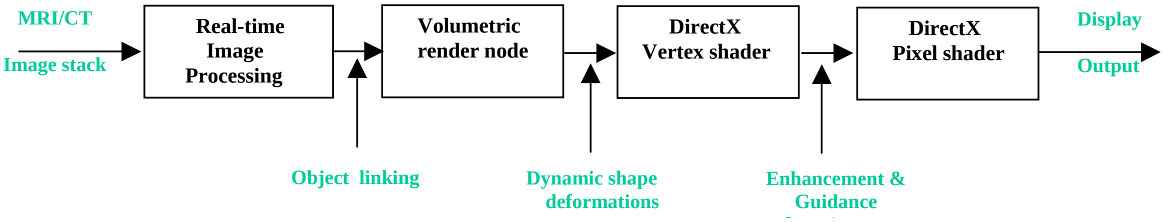

Figure 4*: Dynamic shading and real-time image processing on the parallel GPU hardware to support anatomical guidance and in-volume structure visualization.*

Proceeding from left to right the functionality of each module is as follows: The volumetric data set (MRI/CT) at the input of the visualization pipeline is defined as a stack of dynamic images (gray scale, color and/or label maps attached). In the first stage real-time image processing operators are used to automatically process all slices within a given volumetric object and implement the required preprocessing steps. These steps comprise of removing unwanted parts, image operators for noise reduction, contrast enhancement or substructure masking. Using the processed slices, a volumetric object is created and linked to an anatomical skeletal structure by an *object linking* mechanism. The resulting volume object then enters the rendering pipeline where it first passes through a vertex *shader* and subsequently in a pixel *shader*. The vertex *shader* implements *local shape deformations* (not used in this paper) while the *pixel shader* implements the algorithms required to highlight different internal structures inside the volume to help guide the US scanning process. The *pixel shader* can use gray scale images or color scheme for best viewing. More specifically, it can be programmed to color the pixels within a volume according to local image density, gray scale color or other algorithmically extractable features and its operation maybe directly controlled via label maps as well. The unit thereby offers an extended set of parameters to visualize anatomically important information a medic needs to scan. The key advantage of this approach is that it delivers high performance and constant visualization speed even on sub-optimal, portable computers as we will later show in Section 5.

The visualization algorithm implemented on the GPU relies on label maps that store segmented structures and allow the shader code to display them only when necessary. More specifically *label maps* are the result of complex segmentation algorithms used in *Slicer*. They identify anatomical structures by assigning them to membership sets. For each anatomical volume its corresponding label map is another volume used to store the resulting output of segmentation method whereas each *voxel* is labeled according to its tissue type. Therefore to use label maps to support 3D guidance we first needed to address the performance challenges associated with storing multiple large volumes in the memory of the graphics card, next to allow the architecture to address multiple 3D texture volumes and finally to use this information to govern the operation of the massively parallel shader codes for visualization. Figure 5 shows typical label maps referring to the test data set. In this torso section there are 478 different segmented anatomical structures with label values ranging from 0 to 2232. This information was recaptured and encoded in the gray intensity values of Figure 5 and fed as input to the pixel shader along with the original color volume.

Figure 5*: Example slices from the label map volume generated to represent the 478 different anatomical structures in the torso data set and used in our visualization algorithms.*

Using the label map data sets show in Figure 5, any segment of the color volume may be visualized at constant speed. As an example Figure 6 shows the output of the algorithms for structures of different complexity using two different shader settings. The upper image shows segmented structures just under the skin surface with translucent skin layer for better visibility. The image below shows another combination of the anatomical structures of interest located inside the body.

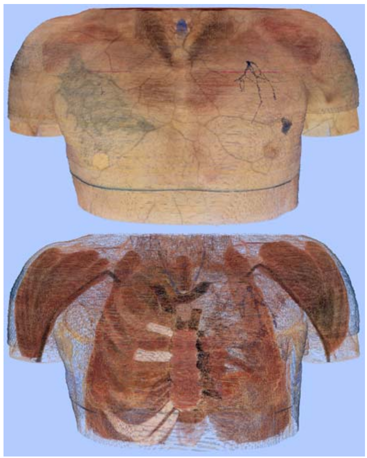

Figure 6*: High-speed visualization of segmented anatomical structures within a volume driven by label maps at the shader level.*

Figure 7 displays the fragment code used in the system. The operation of the pixel shader which takes a 3D texture (*volumeTex*) as input is governed by a number of parameters, called constants all stored in registers (*baseRange, highlightColor, highlightDelta, massColor in c0 – c3).* These constants are four element vectors with x,y,z,w coordinates. For each voxel's output (*color*) the red, green, blue and transparency or alpha (*r,g,b, and a*) values are computed via a simple algorithm and uploaded to the graphics card. To demonstrate the visualization capabilities of the above pipeline, Figure 8 shows examples of 3D visualization of anatomical structures of the same data set as in Figure 2. In the upper two rows the same view is shown with different shader parameters, while in the rows below the torso from different viewing angles for a constant set of visualization parameters is shown.

```
  struct vertout
  {
     float4 Pos : POSITION;
     float4 Col  : COLOR0;
     float4 uvw : TEXCOORD0;
  };


sampler3D    volumeTex      : register( s0);
float4         baseRange     : register( c0);
float4         highlightColor : register( c1);
float4         highlightDelta : register( c2);
float4         massaColor     : register( c3);

float4 main( vertout IN) : COLOR
{
  float4 color = (float4)1;
  float4 texCol                    = tex3D(volumeTex, IN.uvw.xyz);

  color.rgb    = texCol.rgb;
  if ((texCol.a > baseRange.x) && (texCol.a < baseRange.y))
    color.a  = highlightColor.w;
  else
    color.a = 0;

  return color;
}
```

Figure 7*: Pixel shader code used for in-volume segmentation and visualization.*

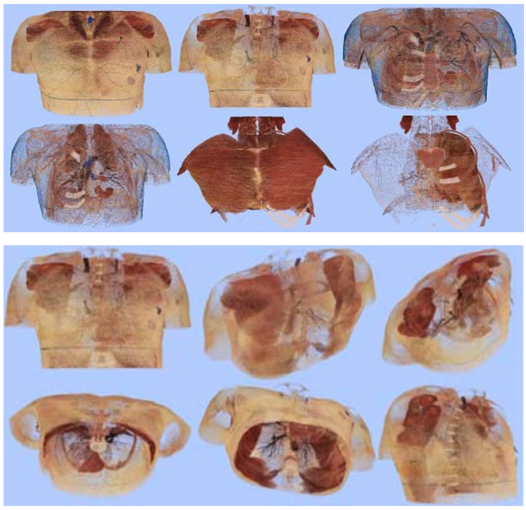

Figure 8*: Visualizing complex anatomical structures.*

When using this shader various constant settings (*c0 through c3*) refer to different visibility and transparency values for each pixel. The user interface of the system than allows for changing these constants to best fit the needs for showing internal structures. An example of this is shown in Figure 9. To take this concept even further these shader parameters can be dynamically uploaded and changed in real-time to best suite the needs of guidance and visualization. Different sets of parameters may be grouped together, interpolation schemed allow medics to interpolate and navigate these settings by a simple interface (*Disc Controller).*

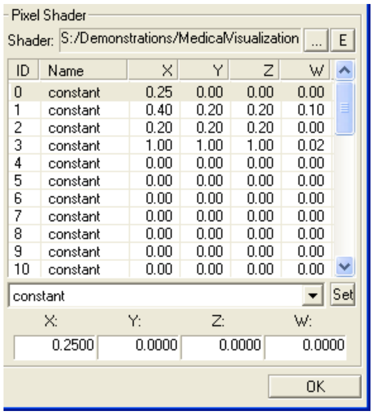

Shader: S:/Demonstrations/MedicalVisualization

ID	Name	X	Y	Z	W
0	constant	0.25	0.00	0.00	0.00
1	constant	0.40	0.20	0.20	0.10
2	constant	0.20	0.20	0.20	0.00
3	constant	1.00	1.00	1.00	0.02
4	constant	0.00	0.00	0.00	0.00
5	constant	0.00	0.00	0.00	0.00
6	constant	0.00	0.00	0.00	0.00
7	constant	0.00	0.00	0.00	0.00
8	constant	0.00	0.00	0.00	0.00
9	constant	0.00	0.00	0.00	0.00
10	constant	0.00	0.00	0.00	0.00

constant

X:	Y:	Z:	W:
0.2500	0.0000	0.0000	0.0000

Figure 9*: Pixel shader user interface and settings of constants for visualizing vascular structures.*

5. Evaluation

As stated above one of the key advantages of our proposed approach is to use volumetric representation in combination with label maps, a technique that delivers high performance and constant visualization speed even on portable computers. To test this assumption we measured the overall performance of the rendering algorithm on an Apple macMini computer (1.66GHz Intel Core Duo, 2GB 667 DDR2 SDRAM) we used in our "backpack" ultrasound guidance system.

The speed tests comprised of two steps. First we used the polygonal models (see Figure 2) of increasing complexity in terms of polygon counts and recorded the overall rendering speed in frames-per-second (fps). Real-time performance requires a minimum of 15 fps update rate for the ultrasound operator to see smooth motion and perceive the system's reaction time as seamless. As shown in the figure below even with relatively small 3D polygonal models the Apple macMini did not reach this performance level. This is largely due to the relatively slow performance of the built-in graphics chip that is quite significantly slower than a high-end graphics card would be. In the second test the same anatomical structures are visualized using the label-map volume algorithm to compare the two methodologies. Our findings are summarized in Figure 10.

The graph shows the *rendering speed* (i.e. how fast the computer reacts and updates images on the screen during anatomical guidance) as a function of scene *complexity* (i.e. how many visual elements are shown). The blue line is the performance curve for 3D models. A typical guidance system shows a few major landmarks, bones, lungs, liver, kidney, vascular structure, etc. The blue curve demonstrates that more structures are shown the slower frame rate becomes, eventually reaching only two frames per seconds (fps). On the other hand the *labelmap-based shading algorithm* delivers much higher speeds (30fps) on the same hardware and a constant performance even when hundreds of internal structures are shown.

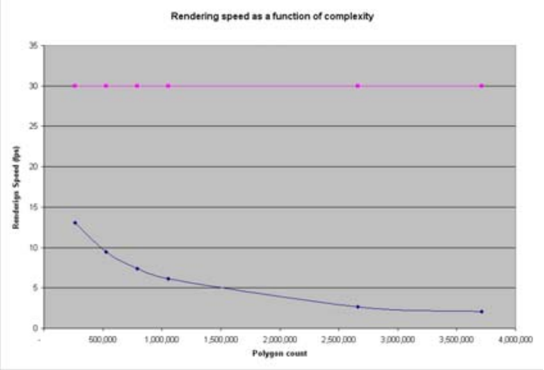

Figure 10: *Comparison of Rendering speed as a function of the complexity of the anatomical structures (see text).*

6. Conclusion

In this paper we introduced a novel methodology to visualize large numbers of complex anatomical structures for 3D guidance purposes. Our system replaces polygonal models with colored volumetric representation combined with pre-segmented label maps. To test our approach we used a detailed 3D models segmented from the torso section of the NIH Visible Male data set and experimentally showed that our algorithm delivers high performance and constant speed visualization outperforming traditional polygon-based methods on the same portable computer hardware.

The visualization method described herein was successfully integrated into an advanced ultrasound guidance framework and used to aid operators to find anatomical structures and obtain high quality ultrasound image sequences. Future work involves constructing reference data sets for the entire volume data made available and generalization of our algorithms to handle multiple volumes in parallel.

References

[1] J. Elsmere, J. Stoll, D. Rattner, D. Brooks, R. Kane, W. Wells III, R. Kikinis, K. Vosburgh (2003), "A Navigation System for Augmenting Laparoscopic Ultrasound", 6th Int. Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'2003), LNCS 2879, pp. 184-191,Montreal, Canada.

[2] B. Takacs, K. G. Vosburgh (2008), "Compact Anatomically Guided Ultrasound for Casualty Care", ACHI08, Sainte Luce, Martinique (under review).

[3] National Library of Medicine, Visible Human Data Set, http://www.nlm.nih.gov/research/visible/getting_data.html

[4] 3D Slicer, www.slicer.org