

A Novel Method for Visualization of Entire Coronary Arterial Tree

THOMAS WISCHGOLL,¹ JOERG MEYER,² BENJAMIN KAIMOVITZ,³ YORAM LANIR,³ and GHASSAN S. KASSAB^{4,5,6}

¹Department of Computer Science and Engineering, Wright State University, Dayton, OH, USA; ²Department of Biomedical Engineering, University of California, Irvine, CA, USA; ³Department of Biomedical Engineering, Israel Institute of Technology, Haifa, Israel; ⁴Department of Biomedical Engineering, Indiana-Purdue University, Indianapolis, IN 46202, USA; ⁵Department of Surgery, Indiana-Purdue University, Indianapolis, IN 46202, USA; and ⁶Department of Cellular and Integrative Physiology, Indiana-Purdue University, Indianapolis, IN 46202, USA

(Received 6 April 2006; accepted 1 February 2007)

Abstract—The complexity of the coronary circulation especially in the deep layers largely evades experimental investigations. Hence, virtual/computational models depicting structure-function relation of the entire coronary vasculature including the deep layer are imperative. In order to interpret such anatomically based models, fast and efficient visualization algorithms are essential. The complexity of such models, which include vessels from the large proximal coronary arteries and veins down to the capillary level (3 orders of magnitude difference in diameter), is a challenging visualization problem since the resulting geometrical representation consists of millions of vessel segments. In this study, a novel method for rendering the entire porcine coronary arterial tree down to the first segments of capillaries interactively is described which employs geometry reduction and occlusion culling techniques. Due to the tree-shaped nature of the vasculature, these techniques exploit the geometrical topology of the object to achieve a faster rendering speed while still handling the full complexity of the data. We found a significant increase in performance combined with a more accurate, gap-less representation of the vessel segments resulting in a more interactive visualization and analysis tool for the entire coronary arterial tree. The proposed techniques can also be applied to similar data structures, such as neuronal trees, airway structures, bile ducts, and other tree-like structures. The utility and future applications of the proposed algorithms are explored.

Keywords—Coronary vasculature, Geometry reduction, Large-scale visualization, Occlusion culling, Tree-shaped data set.

INTRODUCTION

To understand such a complex system as the coronary circulation, it is essential to employ anatomically based mathematical models that integrate the physical and biological interactions. It is important for these

virtual models to include high detail at the microvasculature (including capillary vessels) as well as on a macroscopic scale (epicardial vessels) in order to integrate the entire coronary vasculature. A visual representation of the anatomical model should include the various parameters of the model. For example, diameters and lengths and their relative changes throughout the vasculature should be visualized for every vessel segment. The visual representation should enable a user to better analyze the parameters of the data set compared to tabular data. In addition, further information should be accessible to the user by selecting a vessel segment and displaying information, such as vessel volume and surface area. The system should also allow the user to edit the individual vessel segments and change their radii or location. Obviously, representing the entire geometry of the vasculature results in a huge set of geometrical data. Ideally, the visualization should be interactive; i.e., the rendering algorithm has to output at least several frames per second (fps).

Rendering such a large-scale model is quite challenging for currently available computing hardware since commodity graphics cards are presently not able to display this amount of information interactively. For the complete coronary arterial model, a total of 6 giga-byte (GB) of geometric information is needed to be transferred from main memory to the graphics card, which presents a limit for interactive rendering. Furthermore, most desktop computers are not capable of handling this amount of data due to insufficient main memory. Hence, the size of such a large-scale anatomical model is prohibitive for rendering on desktop computers without employing out-of-core techniques.

The objective of this study is to develop a visualization method for a view-dependent, interactive decimation of massive tree-shaped data sets. The proposed

Address correspondence to Ghassan S. Kassab, Department of Biomedical Engineering, Indiana-Purdue University, Indianapolis, IN, 46202, USA. Electronic mail: gkassab@iupui.edu

86 approach will combine a spatial data structure and
 87 occlusion queries to reduce the number of triangles
 88 necessary to render tree-shaped data sets that exceed
 89 the memory present in the computer system. The
 90 topology of tree-shaped data sets is exploited in order
 91 to reduce the complexity of the triangle mesh coher-
 92 ently. The proposed software system makes use of
 93 recent improvements in graphics hardware and
 94 employs hardware occlusion queries that allow a faster
 95 and more precise occlusion test as compared to soft-
 96 ware-based approaches. The techniques described in
 97 this article can be easily applied to data extracted from
 98 any tree-like structures.

99 METHODS

100 *Anatomically Based Model*

101 Recently, Kaimovitz *et al.*¹⁵ developed a three-
 102 dimensional (3-D) geometric model of the entire cor-
 103 onary arterial tree (right coronary artery, RCA; left
 104 anterior descending artery, LAD; and left circumflex,
 105 LCx arterial tree) based on Kassab *et al.*'s coronary
 106 morphometric data base.¹⁶ The model spans the entire
 107 coronary arterial tree down to the capillary vessels in a
 108 prolate spheroid model of the heart and encompasses
 109 about 10 million segments. The 3-D tree structure was
 110 reconstructed initially in rectangular slab geometry by
 111 means of global geometrical optimization using a
 112 parallel Simulated Annealing (SA) algorithm. The SA
 113 optimization was subject to a global boundary avoid-
 114 ance constraint and local constraints at bifurcations
 115 prescribed by previously measured data on branching
 116 asymmetry in the coronary arterial tree.³⁸ Subse-
 117 quently, the reconstructed tree was mapped onto the
 118 prolate spheroidal geometry of the heart. The trans-
 119 formation was made through least squares minimiza-
 120 tion of the deformation in segment lengths as well as
 121 their angular characteristics.

122 *Rendering of Massive Tree-Like Structures*

123 In the previous publication,¹⁵ vessel segments were
 124 visualized using standard cylinders. Since consecutive
 125 vessel segments do not necessarily form 180 degree
 126 angles, these result in visible gaps at the point of
 127 transition. To avoid these gaps, the proposed system
 128 represents vessel segments as conic cylinders with
 129 rotated ends, which are not necessarily orthogonal to
 130 the cylindrical axis. In this way, a smooth transition
 131 from one segment to the daughter segment(s) can be
 132 achieved, thus avoiding any gaps. The individual conic
 133 cylinders are pieced together using triangles that are
 134 fitted in such a way that an optimal, gap-less approx-
 135 imation is achieved. This results in an accurate visual

representation of the entire vascular structure as 136
 defined by the data set. 137

Since several triangles are needed to represent a 138
 single conic cylinder, rendering a vascular structure 139
 which consists of 10 million vessel segments requires 140
 about 220 million triangles to achieve a sufficiently 141
 accurate approximation. This in return results in 142
 geometry data that amounts to several GB in size 143
 which exceeds the main memory of common desktop 144
 computers. In addition, transferring this amount of 145
 data to the graphics hardware and processing this 146
 information overwhelms both the bus system (usually 147
 advanced graphics port, AGP, or PCI Express) as well 148
 as the graphics hardware. Consequently, techniques 149
 are needed that allow the system to handle data sets 150
 that exceed the amount of main memory present in the 151
 computer as well as reduce the number of triangles to 152
 generate the visualization. 153

Hence, the proposed software system deploys out- 154
 of-core techniques which store the entire geometry 155
 data on the hard drive only. During the rendering 156
 process, only parts of the data are transferred to the 157
 main memory. Once these parts are processed, the 158
 system automatically removes these parts and loads 159
 the next ones for further processing. In this way, the 160
 geometry data is loaded in a streaming fashion from 161
 the hard drive and then transferred to the graphics 162
 hardware for visualization. 163

In addition, the proposed system reduces the number 164
 of triangles using view-dependent geometry reduction, 165
 backface-culling, and occlusion-based reduction. View- 166
 dependent geometry reduction automatically reduces 167
 the amount of detail that is used for representing the 168
 vessel segments based on the distance to the viewer. 169
 Hence, the vessel segments that appear far away are 170
 drawn with less detail (using a lower number of trian- 171
 gles per vessel segment) while the ones in the front are 172
 shown in full detail. Since usually only half of a conic 173
 cylinder is visible at a time, only the visible half needs 174
 to be processed in order to generate the visualization. 175
 Accordingly, the number of triangles can be reduced 176
 significantly by removing those triangles of a conic 177
 cylinder that face away from the viewer. 178

Similarly, occlusion-based reduction removes those 179
 conic cylinders that represent vessel segments which 180
 are obstructed by several other vessel segments and 181
 therefore invisible from the current location of the 182
 viewer. Since these vessel segments are not visible, they 183
 can be eliminated without changing the visualization. 184
 This reduces the number of triangles that need to be 185
 transferred to the graphics hardware. Note that all 186
 these techniques for reducing the number of triangles 187
 are view-dependent; i.e., whenever the location of 188
 the viewer changes these need to be recomputed to 189
 ensure that only those triangles are removed that 190

191 minimally contribute to the current visualization. As a
 192 consequence, all computations required need to be
 193 implanted very efficiently. For example, better effi-
 194 ciency can be achieved by grouping vessel segments
 195 that are close together and then applying the geometry
 196 reduction techniques to the entire group. This reduces
 197 the computation effort for visibility tests significantly.
 198 Grouping of vessel segments can be achieved by, for
 199 example, sub-dividing the bounding box of the entire
 200 vascular structure into equal sub-areas. A detailed
 201 description, including implementation details, can be
 202 found in the appendix.

203 RESULTS

204 *Tree Rendering*

205 Figure 1 shows the complete vascular model
 206 including all three major branches. The left image
 207 depicts an overview, while the right displays a close-up
 208 view of the marked region. Even after zooming into the
 209 model, there is still an enormous amount of detail
 210 which underscores the complexity of the generated
 211 model. The geometric model exceeds the main memory
 212 of most desktop computers. In addition, the geometry
 213 data results in slow performance due to the enormity
 214 of information that needs to be processed to compute a
 215 single projected image. Without applying out-of-core
 216 methods, only one branch of the vasculature, for
 217 example the LCx as depicted in Fig. 2 (a), can be
 218 visualized since it is significantly smaller in size
 219 (1.8 million vessel segments).

220 *Geometry Reduction*

221 The geometry reduction techniques applied to the
 222 data significantly reduced the number of triangles. For

example, Fig. 2(b) shows the results for the view-
 dependent geometry reduction where 20% less triangles
 were required for generating the image. Figure 3 shows
 an example of hardware occlusion queries applied to
 the LCx data set. A reduction of 56% was achieved.
 The vessel segments in those areas that were identified
 as occluded are colored in red and drawn at the lowest
 level of detail. As can be seen in the figure, only
 those parts that are far away from the observer and
 obstructed to a large extent by other vessel segments
 are displayed using a lower level of detail (red).

234 *Performance*

235 For performance testing, two different systems were
 236 used. The first one was a Pentium4 2.6 giga-Hertz
 (GHz) central processing unit (CPU) equipped with
 2 GB of main memory and an AGP version of an
 Nvidia GeForce fx5200 graphics card. The second one
 was equipped with two AMD64 Opteron 246 2.0 GHz
 processors and 1 GB of main memory. This system
 used the PCI Express version of an Nvidia Quadro FX
 4400.

244 As one of the reference data sets, the LCx coronary
 245 artery data was visualized. For displaying the data
 set in full detail (considering all vessel segments at
 full resolution), the geometry consisted of 25 million
 triangles. On the first test system without any
 enhancements, the data set could be rendered at an
 average frame-rate of 0.5 fps. Figure 4(a) provides a
 comparison between the frame rates obtained for the
 full data set and for the instantly reduced data set using
 occlusion culling. The frame rates over time are
 depicted for three different rendering methods using
 the same data set, and are shown while navigating
 through the model. In full level-of-detail mode, the




FIGURE 1. Complete representation of the vasculature of a heart and close-up view depicting the large amount of detail in the model.

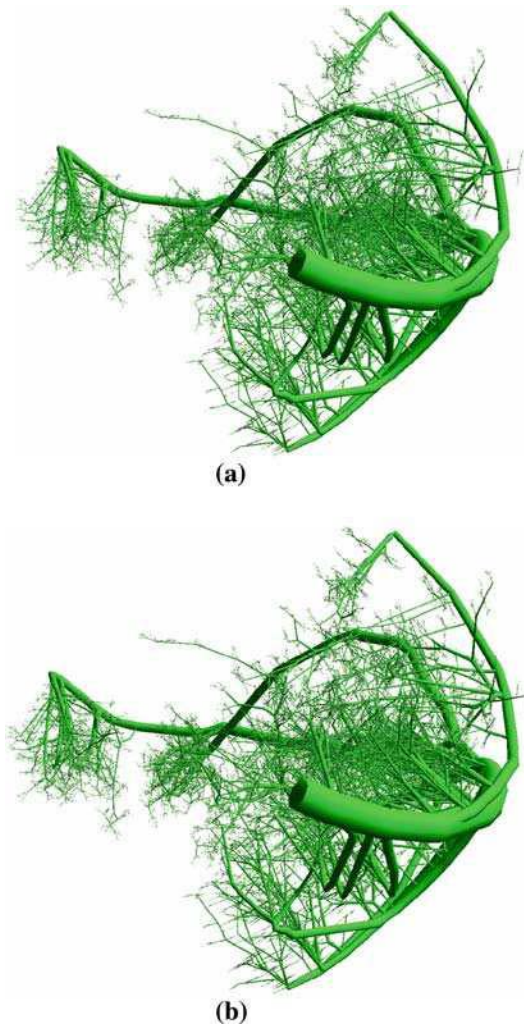


FIGURE 2. Rendering of the geometry of the left circumflex coronary artery (LCx) data set at full resolution (a) and LCx branch rendered with view-dependent geometry reduction enabled (20% reduction) (b).

257 graphics hardware could not be used to its full extent
 258 when the first computer system was used. By enabling
 259 occlusion culling combined with view dependent
 260 geometry reduction, the number of triangles that need
 261 to be displayed for each frame was trimmed down to
 262 an average of 11 million triangles according to
 263 Fig. 4(b). Consequently, the frame rate increased to an
 264 average of 2.1 fps due to the reduced number of tri-
 265 angles. Since the performance increased by a factor of
 266 four while cutting the number of triangles in half only,
 267 the saturation of the AGP bus and CPU of the test
 268 system is improved resulting in more efficient usage of
 269 the graphics hardware.

270 For the occlusion test, the number of sub-areas used
 271 should not be too large since the more time is spent on
 272 occlusion culling the less time is available for the actual
 273 rendering of the image. In this case, a heuristically

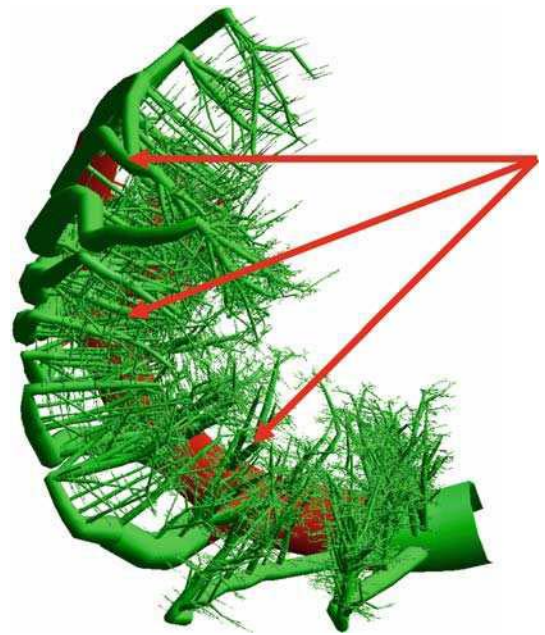


FIGURE 3. Close up of the LCx branch rendered with hardware occlusion culling enabled (56% reduction). Areas rendered with reduced resolution are shown in red as marked by arrows.

determined equidistant scheme of $10 \times 10 \times 10$ sub- 274
 volumes was used for the performance tests. Due to the 275
 fact that the number of triangles that need to be ren- 276
 dered for each frame could be reduced to 11 million 277
 triangles, the utilization of the AGP bus is improved. 278
 As a result, the rendering system was able to render 279
 about 23.1 million triangles per second. Therefore, to 280
 achieve a significantly better rendering performance 281
 using the graphics hardware, we incur only a minor 282
 performance loss for conducting the occlusion tests. 283

The performance of the backface culling imple- 284
 mented in the system was tested on the RCA data set 285
 (consists of 4.3 million vessel segments) which was 286
 represented by 77 million triangles. The data set was 287
 rendered on the system equipped with an Nvidia 288
 GeForce fx5200 graphics card as previously described 289
 and utilized the implemented out-of-core technique. 290
 Figure 5(a) shows the number of triangles used during 291
 rendering. The backface culling was done in a con- 292
 servative way where only about one third of the tri- 293
 angles were removed. In this way, only invisible 294
 triangles are removed. This is especially necessary 295
 when rendering vessel segments that are represented by 296
 a very low number of triangles. For example, for a 297
 vessel segment represented by eight triangles, six of 298
 these can be seen in the worst case. According to 299
 Fig. 5(a), about 50 million triangles were required for 300
 rendering after removing those triangles that face away 301
 from the view point. This then increased the rendering 302
 rate accordingly as can be seen in Fig 5(b). Originally, 303

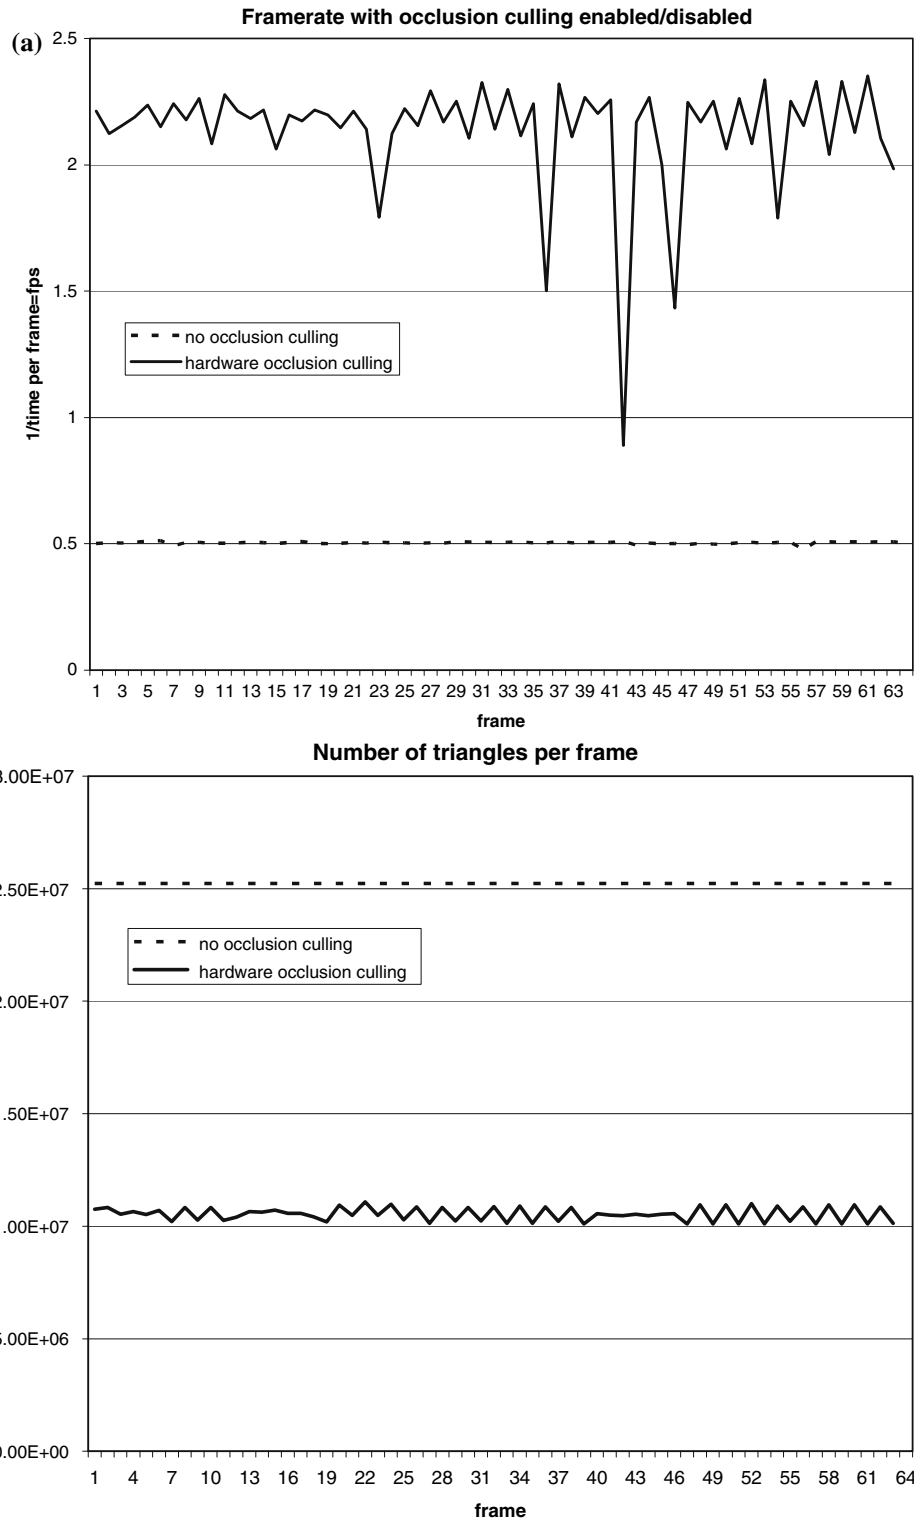


FIGURE 4. Resulting frames per second (a) and number of triangles used (b) when displaying an overview rendering of the left circumflex coronary artery (LCx) branch on a desktop PC equipped with an Nvidia GeForce fx5200 graphics card.

304 a rendering speed of an average of 0.23 fps was
 305 achieved. After removing backfacing triangles, the
 306 data set is rendered at 0.33 fps, a 39% improvement in

performance. During rendering, the software system
 had a memory footprint of 1.6 GB mainly used for
 caching most of the data set.

307
 308
 309




FIGURE 5. Number of triangles used (a) and frames per second (b) when rendering an overview of the right coronary artery (RCA) branch on a desktop PC equipped with an Nvidia GeForce fx5200 graphics card with and without backface culling.

310 The out-of-core technique was tested using all three
 311 branches of the coronary arterial tree model. The
 312 geometry of this model is represented by 220 million
 313 triangles to represent all 10 million vessel segments
 314 resulting in 6 GB of geometry information. Rendering
 315 was performed both on the system equipped with an
 316 Nvidia GeForce fx5200 graphics card as well as an
 317 Nvidia Quadro FX 4400. According to Fig. 6(a),
 318 rendering a single frame of this data set on the first
 319 system took about 3 min and 52 s. Using the second
 320 system, rendering the full data set took only about 62 s
 321 as can be seen in Fig. 6(b). Obviously, the implemented
 322 visualization system benefits from the faster graphics
 323 hardware and the 64bit architecture available in the

test system. Due to the out-of-core visualization, the
 full model could be rendered using less than 64 MB of
 main memory as observed via the Windows task
 manager. As pointed out previously, the out-of-core
 approach exploiting standard memory mapping tech-
 niques benefits from the caching capabilities of the
 operating system especially well when the user zooms to
 a certain area so that the geometry required for ren-
 dering this part fits into main memory. Figure 7 shows
 the performance of the rendering for such a case. An
 average of 1.9 fps was achieved. The system was able to
 render the data at more than two fps. Only in those
 cases where it is required to load the geometry from a
 sub-area that was not displayed before, the system's

Visualization of Entire Coronary Arterial Tree




FIGURE 6. Resulting frames per second for rendering an overview of the complete geometric model of the coronary artery tree on two different desktop PCs, one PC equipped with an Nvidia GeForce fx5200 graphics card (a), the other one equipped with an Nvidia Quadro FX 4400 graphics card (b).

338 performance dropped. Overall, the system was capable
 339 of rendering even huge data sets as can be seen from
 340 these examples and allows visualization at respectable
 341 frame rates even with small main memory computers.

DISCUSSIONS

343 This study provides a visualization method for a
 344 quantitative anatomical model of tree structures (such
 345 as coronary arterial trees) that can be used, for instance,
 346 to model the temporal and spatial distribution of blood
 347 flow in the heart. In this study, the described simulated
 348 data set was visualized. It is possible, however, to use

the system for other types of data sets, such as microCT
 scanned specimens^{20,26} or data retrieved by an imaging
 cryomicrotome.³² The visualization features of the
 software will serve as an educational tool as well as for
 data interpretation. It is expected that the software will
 become a valuable tool for cardiologists, physiologists
 and students. The details are discussed below.

Visualization

When rendering such a large model, there are three
 factors that can limit performance. First, the amount
 of information can saturate the bus system so that the
 amount of data cannot be transferred fast enough to

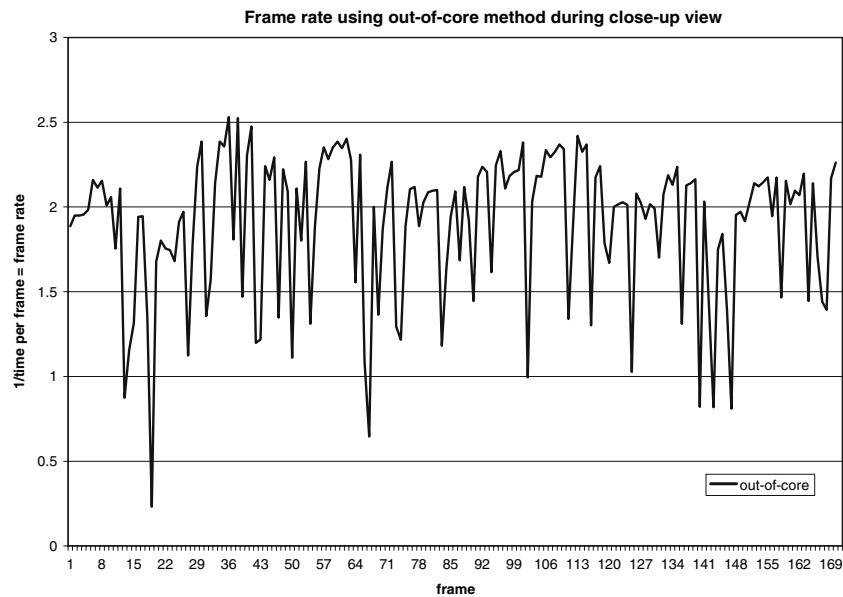


FIGURE 7. Resulting frames per second for rendering the complete geometric model of the coronary artery tree on a desktop PC equipped with an Nvidia GeForce fx5200 graphics card during close-up view.

361 the graphic card. Even though the PCI Express (PCIe)
 362 bus at 2.5 Gbit/s is faster than the advanced graphics
 363 port (AGP) with transfer rates of up to 2 Gbit/s, both
 364 bus systems can be saturated due to the large size of the
 365 model. Second, the capabilities of the graphic card
 366 itself may not be sufficient to render the data fast
 367 enough. Third, there may not be enough main memory
 368 available to store the entire data set. Common 32-bit
 369 desktop computers can only be equipped with up to
 370 4 GB of main memory which is not enough for storing
 371 the entire geometric representation. If additional
 372 information is to be displayed, such as colors for the
 373 vessel segments to superimpose pressure or flow, this
 374 increases even further. However, the memory require-
 375 ments for storing color information are significantly
 376 less than those for the geometry (about 1%) so that it
 377 would not reduce the performance significantly.

378 Sixty-four bit computers are becoming available,
 379 which can be equipped with up to 16 GB of main
 380 memory. BIOS and driver limitations, however, often
 381 still limit the amount of memory to 3 GB. Whenever
 382 the available memory is not sufficient for storing the
 383 entire geometry, the hard disk needs to be used as
 384 secondary storage which results in significantly slower
 385 data access.

386 In order to alleviate these limiting factors, the
 387 number of triangles employed to display the model
 388 needs to be reduced resulting in less data that needs to
 389 be processed. Suitable techniques include geometry
 390 reduction techniques which implies that the geometric
 391 model is displayed with a reduced number of triangles
 392 in those areas that are far from the view point. In

393 addition to a view-dependent level-of-detail represen-
 394 tation based on distance, occlusion-based methods can
 395 be used. These techniques allow the identification of
 396 areas in the model that contribute little to the final
 397 image. This is usually due to the fact that the segments
 398 located in those areas are obscured by other segments
 399 which are positioned more closely to the view point
 400 when projected onto the viewing plane. These occluded
 401 areas of the model can then be either eliminated or
 402 rendered using a lower resolution and consequently
 403 reduce the overall number of triangles.

404 Tree-shaped geometric structures have certain
 405 unique properties that render most traditional occlu-
 406 sion culling algorithms inefficient. For instance, when
 407 rendering architectural models or iso-surface repre-
 408 sentations of objects, occlusion frequently occurs. As
 409 an example, if the camera is located inside a room with
 410 no windows of an architectural model, the entire out-
 411 side world is not visible, thus occluded. This is not
 412 likely to happen in tree-shaped data sets because the
 413 scene consists predominantly of relatively skinny ele-
 414 ments which make partial occlusion much more likely
 415 than complete occlusion. Consequently, occlusion
 416 query techniques can only be used as a measure for
 417 visibility indicating the extent to which the precision of
 418 the model can be reduced without changing the visual
 419 appearance very much. Due to the limited occlusion
 420 within such a vascular tree structure, occlusion culling
 421 should not be used as a method for completely elimi-
 422 nating parts of the scene because certain areas may still
 423 be partially visible in most cases. Therefore, the present
 424 system uses OpenGL occlusion queries over the

425 GL_HP-occlusion-test which allows the system to
426 determine the amount of occlusion.

427 *Comparison with Other Studies*

428 The resulting geometrical representation of the
429 simulated vascular tree consists of millions of vessels.
430 Due to the tremendous complexity of this model, the
431 images shown in the original article¹⁵ were generated
432 using the POVray ray tracer. Those only include ves-
433 sels down to order five to prevent POVray from being
434 overwhelmed by the complexity of the model. It took
435 about two minutes to generate an image of the reduced
436 vasculature using POVray. In this article, however, all
437 vessels down to the capillaries are shown and less time
438 is needed for creating an image.

439 The proposed method is based on a real-time, view-
440 dependent simplification of complex models. Several
441 publications exist that employ similar methods.
442 Progressive meshes as introduced by Hoppe¹³ were
443 designed to obtain increasingly coarser representations
444 of a mesh by applying edge collapse operations.
445 Applying this method, a level-of-detail description of
446 the model is derived. In one of his later publications,
447 Hoppe¹⁴ describes efficient data structures and algo-
448 rithms for implementing progressive meshes. Xia
449 et al.³⁴ defined the notion of a *merge tree* that stores
450 the edge collapse operations in a hierarchical manner
451 to create a continuous-resolution representation of an
452 object. A similar approach was proposed by El-Sana
453 et al.⁵ where a binary view-dependence tree is created
454 containing general vertex-pair collapses. This tree can
455 then be used to generate the required triangles for
456 display at run time. Andujar et al.¹ used classical
457 occlusion culling algorithms and computed potentially
458 visible sets (PVS) which consist of those polygons that
459 are likely to be visible. These sets are supersets of the
460 sets of all visible polygons for which the degree of
461 visibility is determined to create view-dependent
462 occlusion culling. Shaffer et al.³¹ developed a pro-
463 gressive mesh simplification algorithm which clusters
464 the vertices using a BSP-tree resulting in an adaptive
465 simplification of the polygonal mesh. Pajarola²³
466 introduces *FastMesh* which defines a hierarchy on half-
467 edges that reduces the storage cost in comparison to
468 vertex hierarchies. El-Sana et al.⁶ combine a view-
469 dependence tree with spatial sub-division techniques to
470 avoid scanning of active nodes that do not contribute
471 to the incremental update of the selected level of detail.

472 Several algorithms for reducing the complexity of a
473 scene using occlusion culling are available both
474 implemented in software and in hardware.³ Greene¹⁰
475 developed an algorithm based on hierarchical tiling
476 that is able to determine whether a convex polygon is
477 inside, outside, or intersecting an image hierarchy.

Bartz et al.² render bounding volumes into a virtual
478 occlusion buffer using OpenGL and read back the
479 results from the graphics hardware to determine
480 occlusion. Since reading back from the OpenGL buffer
481 is slow, an interleaving scheme is applied to speed up
482 read-back. Zhang et al.³⁹ describe hierarchical image-
483 space occlusion maps for visibility culling. The culling
484 algorithm uses an object-space bounding volume
485 hierarchy and can be implemented using graphics
486 hardware. Klosowski et al.¹⁷ propose a visibility cull-
487 ing algorithm based on *Prioritized Layered Projection*
488 (*PLP*) that can be implemented using graphics hard-
489 ware. El-Sana et al.⁶ combine the PLP approach with
490 view-dependency resulting in a view-dependent occlu-
491 sion culling. Yoon et al.³⁶ use a clustering hierarchy
492 for refining the underlying grid to obtain a level-
493 of-detail representation for arbitrary triangle meshes in
494 addition to hardware occlusion culling. Recent efforts
495 show that current hardware improvements and the
496 usage of a *clustered hierarchy of progressive meshes* can
497 improve rendering speed even further.³⁷ However,
498 most of the described methods are not suitable for
499 directly reducing the complexity of a model of tree-like
500 anatomical structures, such as the coronary vascular
501 tree.

502 Different techniques for visualizing vascular struc-
503 tures can be found in the literature. Gerig et al.⁹
504 describe how to derive a skeletal structure from a
505 volumetric image based on hysteresis thresholding
506 and binary thinning. Hahn et al.¹² employ geometrical
507 primitives, such as truncated cones, to visualize vessels
508 inside the human liver. A similar approach has been
509 taken for the rendering method described in this arti-
510 cle. The model is represented by conic cylinders as
511 previously described. Masutani et al.¹⁸ used cylinders
512 aligned to the vessel skeleton to visualize the vascula-
513 ture. Different radii at branchings resulted in discon-
514 tinuities when using this method. Felkel et al.⁸
515 reconstructed liver vessels from center line and radius
516 information to supply an augmented reality tool for
517 surgery. Puig et al.²⁴ developed a system for exploring
518 cerebral blood vessels using a symbolic model with a
519 focus on geometric continuity and on realistic shading.
520 Oeltze et al.^{20,22} use convolution surfaces to obtain a
521 smoother representation of blood vessels extracted
522 from CT or MR data.

523 Deussen et al.⁴ use points and lines to represent
524 complex systems of plants as approximation reducing
525 the overall number of triangles compared to their ori-
526 ginal representation. Gumhold et al.¹¹ use a splatting
527 approach based on ellipsoids for rendering scientific
528 data sets. The advantage of such a glyph-based
529 approach is the potential of deploying the hardware for
530 rendering. Reina et al.²⁵ showed this when rendering
531 molecular visualizations of 500,000 particles at 10 fps.
532

533 In a comparative study, the present system was
 534 compared to an implementation on a high-end visu-
 535 alization server: the Sun Fire V880z visualization ser-
 536 ver. This server is equipped with 16 GB of main
 537 memory allowing the system to store the entire geo-
 538 metric model in main memory. Despite the fact that
 539 this server is geared towards optimal rendering per-
 540 formance for large data sets, the overall performance
 541 was slower compared to the second test system using
 542 an Nvidia Quadro FX 4400. Generating a single pro-
 543 jected image took 90 s on the Sun server while the
 544 second test system using the present system completed
 545 the task after 62 s.

546 *Significance*

547 Virtual models of normal hearts are needed as a
 548 physiologic reference. Pathological states can then be
 549 studied in relation to changes in model parameters that
 550 alter coronary perfusion. With such computational
 551 models, researchers can analyze the effects of different
 552 treatment options (medical and surgical), and ulti-
 553 mately find rational ways to prevent and treat coro-
 554 nary heart disease. Based on detailed anatomically
 555 based models, computational fluid dynamics simula-
 556 tions can yield accurate simulation of blood flow in
 557 health and disease. In order to visualize the present
 558 anatomically based models that may include future
 559 hemodynamic and physiological data, it is essential to
 560 have efficient and fast visualization techniques. The
 561 present study is the first step in that direction.

562 *Conclusions and Future Work*

563 A rendering system has been described which
 564 exploits the tree-shaped topology to increase rendering
 565 performance. Due to the nature of tree-shaped struc-
 566 tures, hierarchical meshes to obtain different levels of
 567 detail can be generated based on the topological
 568 structure of the data; i.e., individual segments can be
 569 clustered as entities. Geometry reduction techniques as
 570 well as occlusion culling enables the system to render
 571 each frame four times faster than the standard method
 572 that displays the full model directly without simplifi-
 573 cation methods. For the LCx data set, the number of
 574 triangles can be reduced in such a way that the amount
 575 of geometric information is small enough to be trans-
 576 ferred to the graphics hardware and fast enough to
 577 utilize the full performance potential of the hardware.

578 Using out-of-core techniques, the full model can be
 579 displayed even on computer systems equipped with
 580 relatively small amount of memory since only 64 MB
 581 are sufficient for the algorithm. With a high-end PC
 582 system, rendering can be even faster using out-of-core
 583 techniques compared to workstations equipped with

much more main memory, such as the Sun Fire V880z 584
 visualization server. 585

In the future, ray tracing or ray casting algorithms 586
 will be applied to the data set to explore if there are any 587
 performance benefits from this approach. Additionally, 588
 GPU based methods that render the conic cylinders 589
 completely on the GPU might increase performance by 590
 reducing the amount of information that needs to be 591
 stored in memory and avoids data transfer on the bus 592
 systems. Using other types of geometry approximations 593
 for the tree segments in combination with hardware 594
 based approaches might yield even better performance, 595
 such as line primitives.³² 596

ACKNOWLEDGMENTS 597

The authors would like to thank Falko Kuester, UC 598
 Irvine, for providing the high-end PC system equipped 599
 with an Nvidia Quadro FX 4400. This work was sup- 600
 ported in part by the National Institute of Health— 601
 National Heart, Lung, and Blood Institute Grant 2 R01 602
 HL055554-06 (GSK), by the National Institute of 603
 Mental Health (NIMH) through a subcontract with the 604
 Center for Neuroscience at the University of Califor- 605
 nia, Davis (5 P20 MH60975), by the National Part- 606
 nership for Advanced Computational Infrastructure 607
 (NPACI), Interaction Environments (IE) Thrust 608
 (10195430 00120410). 609
 610

APPENDIX 611

Rendering of Massive Tree-Like Structures 612

In order to analyze large-scale tree-like structures, 613
 appropriate visualization methods are necessary. 614
 Whenever the geometry data of such a structure exceed 615
 the amount of main memory of the computer, the 616
 application of several techniques to both be able to 617
 handle the data set as well as improve performance are 618
 required. This Appendix provides details about the 619
 data format and explains the different techniques that 620
 were applied to visualize the data on common desktop 621
 computers. 622

Visualizing Tree-Like Structures 623

The structure is given as a sequence of consecutive 624
 segments where one segment can have multiple 625
 daughter vessels (mostly two as bifurcations) as suc- 626
 cessors, forming a tree-shaped structure with a highly 627
 asymmetric branching pattern. Each segment in the 628
 tree is characterized by the coordinates and radii of its 629
 proximal and distal nodes. This data format is similar 630
 to the one provided by commercial software packages, 631

632 such as *Analyze*.^{26,28,29} Since the radii of two consec- 670
 633 utive nodes are not necessarily equal, a conic cylinder 671
 634 is defined based on this data resembling each segment. 672
 635 All conic cylinders together then define a representa- 673
 636 tion of the vasculature as prescribed by the model. For 674
 637 a coronary arterial tree, there exist three major branches 675
 638 in the data set representing the RCA, LAD and 676
 639 LCx arterial tree, respectively. Every branch includes a 677
 640 complete set of linked vessel segments from the mac- 678
 641 rovasculature down to the first capillary bifurcation. 679

642 In order to generate a visual representation of the 680
 643 vascular tree, OpenGL and computer graphics techni- 681
 644 ques are used. In computer graphics, a camera 682
 645 analogy is followed similar to taking a photograph. A 683
 646 virtual camera is placed next to the objects, in this case 684
 647 the arterial tree. The orientation of the camera identi- 685
 648 fies the view direction; while the view direction 686
 649 combined with the location of the camera define the 687
 650 view. Using this definition of a view, all objects are 688
 651 then projected onto a virtual image plane. One can 689
 652 think of this image plane as the film inside the camera 690
 653 that was just defined. This projected image is then 691
 654 displayed on the computer screen. Consequently, the 692
 655 view as defined here identifies exactly what parts of the 693
 656 objects are displayed on the computer screen. 694

657 Figure 8 shows an illustration of this configura- 695
 658 tion. Since current graphics hardware does not sup- 696
 659 port the display of complex objects; simple, more 697
 660 universal primitives have to be used. Due to its uni- 698
 661 versal nature, triangles are the most common primi- 699
 662 tives in computer graphics since every complex object 700
 663 can be approximated using a set of triangles. As can 701
 664 be seen in Fig. 8, the conic cylinders that are used to 702
 665 represent the arterial tree are broken down into a 703
 666 series of triangles as well. The triangles are arranged 704
 667 in such a way that they approximate a conic cylinder; 705
 668 i.e., two of the edges of each triangle run along the 706
 669 main direction of the conic cylinder while one edge is 707

670 parallel to one of the end caps of the cylinders. 671
 672 Hence, circles are computed at the transition between 673
 674 vessel segments in a first step. Since cylindrical 675
 676 arteries are considered, these circles are perfectly 677
 678 round. The circles are connected along each vessel 679
 680 segment using triangles to approximate the conic 681
 682 cylinder. For this, the circles are approximated by a 683
 684 series of points by computing a fixed number of 684
 685 points at equidistant locations along the circle. A set 686
 687 of two points (one on the circle at the one end of the 687
 688 cylinder, the second on the opposite side) are then 688
 689 selected in such a way that they form the closest 689
 690 distance between. This forms the first edge of the first 690
 691 triangle. By connecting the next points following the 691
 692 discretization of the circles in an alternating fashion, 692
 693 triangles are formed that approximate the conic cyl- 693
 694 inders. Figure 8 illustrates this by showing a sample 694
 695 triangle imposed on one of the conic cylinders. 695

696 To increase performance, OpenGL provides 696
 697 so-called triangle strips that require less data to be 697
 698 transferred to the graphics hardware for image display. 698
 699 In this case, instead of specifying all three vertices for 699
 700 all triangles, only the vertices for the first triangle need 700
 701 to be specified completely. For the subsequent triangles, 701
 702 only one vertex is specified and the last two vertices of 702
 703 the previous triangle are re-used, forming a new tri- 703
 704 angle that is directly connected to its predecessor. Since 704
 705 the triangles approximating a conic cylinder (and 705
 706 therefore a vessel segment) are all attached to each 706
 707 other, a triangle strip can be used to reduce the number 707
 708 of vertices that need to be specified. As an additional 708
 709 performance increase, OpenGL vertex arrays are used. 709
 710 Vertex arrays require less function calls and hence can 710
 711 be processed by the hardware more efficiently. To use 711
 712 OpenGL vertex arrays, all vertices for a single triangle 712
 713 strip are stored in a consecutive memory area. This 713
 714 memory area can be passed onto OpenGL in a single 714
 715 function call which results in drawing the entire 715
 716 triangle strip. Hence, an entire vessel segment is 716
 717 represented by one vertex array. 717

718 To achieve a smooth transition between consecutive 718
 719 vessel segments, the circle at the end of each cylinder is 719
 720 not necessarily orthogonal to the cylinder itself. 720
 721 Instead, these circles are created in such a way that the 721
 722 plane in which the circle resides divides the angle 722
 723 between the center-lines of the two consecutive conic 723
 724 cylinders into two equal halves (Fig. 8). Using this 724
 725 approach, the center lines of two subsequent segments 725
 726 can form an angle of up to 180 degrees (reverse 726
 727 direction). However, the most common angles in the 727
 728 coronary arterial tree are much smaller. Rotation of 728
 729 the cylinder ends as previously described does not 729
 730 change the way it is approximated by a triangle strip. 730

731 For the smallest of the three branches (LCx arterial 731
 732 tree), the number of triangles that is required for 732
 733 724




FIGURE 8. Interface between two successive vessel segments shown as conic cylinders which consist of several triangles. Camera and view direction are also shown.

725 visualization of all 1.8-million vessel segments is
 726 25 million triangles. The entire coronary tree consist-
 727 ing of 10 million vessel segments is represented by
 728 220 million triangles based on a discretization of each
 729 conic cylinder using sixteen triangles. The conic cylin-
 730 der representing each segment is approximated by a
 731 single triangle strip which is constructed as follows: the
 732 two ends of the conic cylinder are discretized using a
 733 maximum of eight points. These points are then con-
 734 nected with triangles. Using a point on each end
 735 alternately, this results in a triangle strip. At full res-
 736 olution, this triangle strip consists of 16 triangles ren-
 737 dered by 18 vertices. Due to the size and detail of the
 738 data set, the geometry of the vascular structure
 739 requires representation of the coordinates of the ver-
 740 tices using 32-bit floating-point numbers. Lower-pre-
 741 cision representations result in truncated positions of
 742 the vertices and therefore can change the overall
 743 geometry significantly. For correct illumination of the
 744 geometry, normal vectors are included using another
 745 set of three 32-bit floating point values. These normal
 746 vectors are required for computing the correct reflec-
 747 tion of light on top of the conic cylinders used as
 748 geometric representation of the vascular structure. The
 749 normal vectors computed for every vertex based on the
 750 original geometry (the conic cylinders). This yields
 751 significantly better results and achieves an additional
 752 depth cue and therefore a more realistic image. In this
 753 way, a user is much better able to recognize the 3-D
 754 geometry in less time even from a projected 2-D image.
 755 Overall, for representing 220 million triangles using
 756 vertex coordinates and normal vectors, about 6GB of
 757 memory is required in order for the entire geometric
 758 representation of the vascular structure to be stored.
 759 This entire information needs to be processed for every
 760 projected image that is used as visualization of the
 761 vascular structure.

762 *Geometry Reduction*

763 In order to increase performance, a common
 764 approach is to reduce the amount of geometry infor-
 765 mation that needs to be processed for a single image.
 766 Usually, this is achieved by using a simpler represen-
 767 tation and/or removing parts of the data set that is
 768 either invisible or only visible to a small extent. Com-
 769 pared to arbitrary triangle meshes, tree-shaped data
 770 sets have special topological features that can be taken
 771 advantage of to speed up the visualization. First, the
 772 connectivity between different segments can be used to
 773 simplify the structure by skipping segments. This
 774 results in a simpler representation of the data. Sec-
 775 ondly, the cylindrical shape of the segments can be
 776 used to identify backfacing triangles on a per-segment
 777 basis instead of determining this information for each

778 triangle which removes data invisible due to the pro-
 779 jection. Since the cylinders are rendered as single tri-
 780 angle strips, the connectivity information can also be
 781 exploited when rendering the model; e.g., for backface
 782 culling as described later. Different levels of detail can
 783 be defined based on the precision at which a conic
 784 cylinder is drawn by reducing the number of points for
 785 each delimiting circle of the cylinders.

786 In the current implementation, three levels of detail
 787 are used: a full resolution level where each conic cyl-
 788 nder is represented by 16 triangles, a reduced level
 789 with 8 triangles per cylinder, and a low level of detail
 790 that skips every other segment and renders each
 791 remaining cylinder with 8 triangles. Obviously, the low
 792 level-of-detail mode should only be used in areas far
 793 away from the view point and mostly occluded; i.e.,
 794 covered by a multitude of other vessel segments and
 795 therefore almost invisible. However, since it is almost
 796 completely occluded, a user would need to rotate or
 797 zoom in order to inspect this part of the vasculature.
 798 Once such an area is rotated and therefore more visi-
 799 ble, the system would automatically increase the level
 800 of detail. Similarly, cracks that occur at the transition
 801 between different levels of detail are not noticeable
 802 because these transitions occur sufficiently far from the
 803 view point and only in at least partly occluded areas.

804 In order to decide the resolution for a particular
 805 segment, one could determine the distance between the
 806 current camera position and the segment itself or
 807 determine the number of pixels that would be projected
 808 onto the screen to represent this segment. However,
 809 due to the enormous amount of segments, the com-
 810 putational effort is too costly which would slow down
 811 the rendering speed to several seconds per frame even
 812 for the LCx data set. In fact, computing the distances
 813 between all vessel segments and the camera would take
 814 longer than computing the projected image for the
 815 entire data set without any reduction techniques.

816 To remedy the situation, a spatial data structure is
 817 used. It is essential to the overall performance of the
 818 system that a simple data structure is used which
 819 requires only minimal computation. Hence, a simple
 820 subdivision scheme of the space covered by the data set
 821 is used. This space is equally divided in each dimension
 822 into sub-areas of the same size. Then, only the distance
 823 between the center of this sub-area and the camera
 824 needs to be calculated during the rendering process to
 825 determine the level of detail for the whole sub-area.
 826 Based on a set of thresholds provided by the user, all
 827 segments contained in each sub-area are rendered in
 828 full, reduced, or low level of detail, respectively. These
 829 thresholds describe the distance between the center of
 830 the sub-area and the camera at which the algorithm
 831 will automatically switch to a lower geometric resolu-
 832 tion. Consequently, these thresholds determine the

833 distance to the camera at which the system switches to
 834 a lower resolution. Using this type of geometry
 835 reduction, the number of triangles (e.g., the LCx cor-
 836 onary artery) can be reduced from 25 to about
 837 20 million triangles without introducing noticeable
 838 artifacts.

839 *Backface Culling*

840 To reduce the number of triangles even further, all
 841 triangles that are located at the backside of the conic
 842 cylinders facing away from the camera can be removed
 843 since they are not visible. OpenGL is able to remove
 844 the backfacing triangles but then they still have to be
 845 transmitted to and processed by the graphics card. In
 846 order to avoid transmitting this amount of informa-
 847 tion, these triangles can be identified in software on the
 848 CPU. Of course, a particular vessel segment can be
 849 aligned at virtually any angle to the viewing direction
 850 with respect to the first triangle within the triangle
 851 strip. Therefore, the set of triangles facing backward
 852 can be different for each individual vessel segment.
 853 This implies that the computation has to be done for
 854 each vessel segment individually. Consequently, these
 855 computations have to be carried out in a very efficient
 856 way in order to avoid slowing down the rendering
 857 process. Again, the fact that the topology of a vessel
 858 segment is known can be exploited. Each vessel seg-
 859 ment is represented by a conic cylinder. Consequently,
 860 usually one half of the cylinder is visible, while the
 861 other half is not. The triangles representing the invis-
 862 ible half can be identified using the normal vectors
 863 since these are computed in such a way that they are
 864 always pointing outwards with respect to the conic
 865 cylinder. One approach for identifying those triangles
 866 with normals facing away from the view point is to
 867 check the normals of every triangle individually. This
 868 would represent a computational burden, however,
 869 that would slow down the rendering process. Since the
 870 vessel segments are rendered as triangle strips, both the
 871 visible and the invisible halves are represented by two
 872 sets of consecutive triangles. Thus, in order to identify
 873 the set of back-facing triangles only, the transition
 874 from triangles with normals facing towards the camera
 875 and those pointing away from the camera has to be
 876 found which is significantly less expensive computa-
 877 tionally. Consequently, only the triangles facing the
 878 camera which are visible triangles need to be drawn
 879 resulting in a significant reduction in the number of
 880 triangles that need to be sent to the graphics hardware.

881 *Occlusion-Based Reduction*

882 Another way of reducing the number of triangles
 883 required for a geometric representation of the vascular
 884 structure is to remove triangles that represent conic

885 cylinders which are hidden behind a multitude of other
 886 vessel segments with respect to the current projection
 887 (referred to as occlusion). In a tree-shaped data set,
 888 complete occlusion is not likely to occur since a single
 889 segment does not significantly obstruct the geometry
 890 located behind it. Many of the segments need to be co-
 891 located and packed very densely in a particular area to
 892 occlude other parts of the vascular tree. However,
 893 complete occlusion is not likely to occur. Thus, those
 894 parts of the tree which are detected as (partly) occluded
 895 are still displayed using the lowest level of detail.

896 The present software system employs occlusion
 897 queries implemented in OpenGL 1.5.³⁰ During such a
 898 query, the OpenGL library keeps track of whether the
 899 specified graphical primitives result in pixels actually
 900 drawn to the projected image. In contrast to the
 901 *GL_HP-occlusion_test*, which only returns a binary
 902 true or false result depending on whether pixels were
 903 drawn or not, the occlusion queries defined in OpenGL
 904 1.5 allow the retrieval of the number of fragments
 905 (pixels) that contribute to the current projected image
 906 during the query. Assuming frame coherence (two
 907 consecutive projected images being similar), we can use
 908 these queries to check for occlusion. By drawing a
 909 bounding box of a sub-area, these queries allow the
 910 software to determine how much of a specific sub-area
 911 is visible based on the previous projected image. These
 912 sub-areas are identical with the ones defined by the
 913 spatial data structure used for the previously described
 914 geometry reduction; i.e., an equidistant sub-division
 915 into cubical areas. Based on a user-defined threshold
 916 describing the number of pixels that need to be visible,
 917 the present system can determine the level-of-detail to
 918 be used for the vessel segments contained in this sub-
 919 area. The smaller the threshold the less vessels are
 920 required to occlude a sub-area. Since the vessel seg-
 921 ments are spread over the entire volume relatively
 922 evenly, a more sophisticated sub-division technique
 923 such as binary-space-portioning (BSP) trees or k-d
 924 trees, which sub-divide space recursively at arbitrary
 925 planes instead of using a fixed scheme, would not result
 926 in a significant improvement. Also, a simpler sub-
 927 division scheme allows for faster processing of the
 928 individual sub-elements during testing for occlusion.

929 For each element of the sub-division, a hardware
 930 occlusion query is issued as described above to check
 931 how many fragments pass the depth test; i.e., con-
 932 tribute to the current projected image. This results in
 933 an estimate of how much of the specific sub-area
 934 contributes to the current projected image. To avoid
 935 actual drawing of the bounding boxes, the color mask
 936 is set to zero in OpenGL. Similarly, the OpenGL depth
 937 buffer is marked as read-only to prevent the bounding
 938 boxes from changing the depth values and therefore
 939 occlude each other.

940 Occlusion queries supported on GeForce 3 and
 941 subsequent NVidia GPUs allow many queries to be
 942 performed simultaneously. Therefore, all bounding
 943 boxes that are needed for the occlusion queries are
 944 drawn first. The result of each occlusion query is stored
 945 in the memory of the graphics hardware. To avoid
 946 stalling of the graphics pipeline, the result is read back
 947 only once for all sub-areas after all occlusion queries
 948 are finished. Each occlusion query returns the number
 949 of fragments of the bounding box that would actually
 950 pass the depth test and would have been drawn if the
 951 color mask would not have been set to zero. Conse-
 952 quently, an occlusion query provides a precise measure
 953 of how much of a certain sub-area is occluded. Based
 954 on a user-specified threshold, the rendering system can
 955 then decide whether to draw the vessel segments con-
 956 tained in that specific area at the full level or at a lower
 957 level of detail.

958 *Out-of-Core Rendering*

959 Since the whole data set representing the complete
 960 model of the arterial vascular tree does not fit into the
 961 main memory of regular desktop computers, an out-
 962 of-core method was implemented to support larger
 963 data sets. This technique uses hard drives as main
 964 storage medium, while main memory is only used for
 965 caching data. In this way, only the portion of the data
 966 set which the algorithm is currently using needs to be
 967 present in main memory. The system automatically
 968 updates this portion; i.e., the system loads another
 969 portion of the data set into main memory and removes
 970 another whenever necessary. This enables rendering of
 971 the entire model of the vasculature on current com-
 972 modity hardware. In this approach, the geometry is
 973 determined in a pre-computational step. For each of
 974 the spatial sub-areas that are used during the render-
 975 ing, the triangles needed for displaying all of the vessel
 976 segments contained in this specific sub-area are com-
 977 puted and then stored in a file. After that, the geometry
 978 data can be removed from main memory. Using such a
 979 streaming technique reduces the memory footprint
 980 significantly. In our experiments the memory con-
 981 sumption of the software implementation was less than
 982 64 MB (as observed via the Windows task manager).

983 The geometric representation is pre-computed and
 984 written to a file in form of binary arrays in the same
 985 way it is used by OpenGL when rendered using vertex
 986 arrays. Note that only the full resolution needs to be
 987 stored in the out-of-core file since the lower levels-of-
 988 detail can be derived by masking elements within the
 989 vertex array. Offsets are stored at the beginning of the
 990 binary file as depicted by Fig. 9. This section of the file
 991 allows the system to determine where to find all vertex
 992 arrays within the file for each of the sub-areas used by




FIGURE 9. Out-of-core data file structure reflecting each sub-area within the spatial data structure and their offsets.

the spatial data structure. With this information, the
 exact location within the file can be determined and
 mapped to memory resulting in a pointer to the base
 address of all vertex arrays representing the geometry
 of all vessel segments within a specific sub-area. From
 this starting point, all vertex arrays can be processed as
 if they are stored in memory.

Implementation Details

The visualization system is based on OpenGL.
 Occlusion queries as defined in OpenGL 1.5 are used for
 the occlusion culling. Figure 10 shows sample code to
 use occlusion queries in this context. Similar to display
 lists, vertex arrays, which are significantly faster than
 immediate mode rendering,³⁴ are used for rendering the
 vessel segments. OpenGL allows selecting the vertices of
 a vertex array that are used during the rendering. Thus, a
 lower level of detail can be realized simply by providing
 OpenGL with a subset of indices that represents a lower
 level of detail. This index array can be pre-computed and
 then provided for every segment destined for lower res-
 olution. Obviously, this index array is the same for each
 segment due to the fact that the conic cylinder repre-
 senting a single segment is discretized in the same way
 for each segment. Figure 11 includes the source code
 used for rendering vessel segments using vertex arrays.

```

// create reference ID in OpenGL for this query
GLuint id;
glGenQueries(1, &id);

// disable color and depth mask to ensure that nothing is actually drawn
glColorMask(0, 0, 0, 0);
glDepthMask(0);
glBeginQuery(GL_SAMPLES_PASSED, id);

// draw the test object, i.e. a cube enclosing the current octree element
glBegin(GL_QUADS);
glVertex3dv(leafs[i][j][k].upperleft);
glVertex3d(leafs[i][j][k].lowerright[0],
           leafs[i][j][k].upperleft[1],
           leafs[i][j][k].upperleft[2]);
glVertex3d(leafs[i][j][k].lowerright[0],
           leafs[i][j][k].lowerright[1],
           leafs[i][j][k].upperleft[2]);
glVertex3d(leafs[i][j][k].upperleft[0],
           leafs[i][j][k].lowerright[1],
           leafs[i][j][k].upperleft[2]);
/* ... other faces of the quad ... */

glEnd();

glEndQuery(GL_SAMPLES_PASSED);

// occlusion query finished, store the number of pixels in "result"
GLint result;
glGetQueryObjectiv(id, GL_QUERY_RESULT, &result);

// enable the color and depth mask again
glDepthMask(1);
glColorMask(1, 1, 1, 1);

```

FIGURE 10. Sample code for OpenGL occlusion query to determine the number of fragments (pixels) that would contribute to the current image when drawing all vessel segments contained in a single octree element. The variable result will contain the exact number of pixels that would be drawn for the bounding box of this specific octree element (for the sake of simplicity only the drawing commands for the first face of the bounding box is shown).

```

// out-of-core data for this octree element can be found within the memory
// mapped area, identified by the variable "base"
OutOfCoreData *outofcoredata = (OutOfCoreData *)base;

// determine pointer to the current vessel segment within memory mapped
// area, i.e. skip the header information
char *vertexarray, *normalarray, *pointer =
    (base + sizeof(double) * 9 + sizeof(unsigned int));
OutOfCoreVertexData *outofcorevertexdata;

// draw all vessel segments using OpenGL vertex arrays
for (int i=0; i<outofcoredata->noelements; i++) {
    outofcorevertexdata = (OutOfCoreVertexData *)pointer;

    // determine the pointers to the vertex and normal data needed for
    // rendering within the memory mapped area
    vertexarray = pointer + sizeof(double) * 3;
    normalarray = vertexarray + sizeof(VATYPE) * novertices * 3;

    // declare vertex and normal arrays for OpenGL
    glVertexPointer(3, VATYPE, 0, (VATYPE *)vertexarray);
    glNormalPointer(VATYPE, 0, (VATYPE *)normalarray);

    // draw current vessel segment
    glDrawElements(GL_TRIANGLE_STRIP, 6, GL_UNSIGNED_INT, start);
    trianglecount += 4;

    // advance the data pointer to the next vessel segment
    pointer += sizeof(double) * 3 +
        3 * novertices * (sizeof(VATYPE) + sizeof(VATYPE));
}

```

FIGURE 11. Sample code for drawing all vessel segments within an octree element using OpenGL vertex arrays; the geometry data is retrieved from the memory mapped file. The starting location of the memory mapped area is indicated by the variable "base".

```

// create file handle for input file
filehandle =
    CreateFile (file,
                GENERIC_READ,
                FILE_SHARE_READ,
                NULL,
                OPEN_EXISTING,
                FILE_FLAG_SEQUENTIAL_SCAN || FILE_ATTRIBUTE_READONLY,
                NULL);

if (!filehandle) {
    cerr << "ERROR: cannot open out-of-core file" << endl;
    return;
}

// create a memory mapped file handle
mmaphandle = CreateFileMapping (filehandle,
                                NULL,
                                PAGE_READONLY,
                                0,
                                0,
                                NULL);

// read offsets for identifying the individual octree elements
int i, j, k;
offsets = new unsigned int**[(division+1)];
leafsizes = new unsigned int**[(division+1)];
for (i=0; i<=division; i++) {
    offsets[i] = new unsigned int**[(division+1)];
    leafsizes[i] = new unsigned int**[(division+1)];
    for (j=0; j<=division; j++) {
        offsets[i][j] = new unsigned int[(division+1)];
        leafsizes[i][j] = new unsigned int[(division+1)];
    }
}

// compute the size of the entire header information
generaloffset = sizeof (unsigned int) +
    3 * sizeof (double) +
    sizeof (OctreeArea) +
    (division + 1) * (division + 1) * (division + 1) *
    sizeof (unsigned int) * 2;

// create view for mapping offsets
char *base = (char *)MapViewOfFile (mmaphandle,
                                     FILE_MAP_READ,
                                     0,
                                     0,
                                     generaloffset);

```

FIGURE 12. Memory mapped reading of the header information from the out-of-core file; similarly the geometric information is read from the out-of-core file for every octree element.

1018 The out-of-core rendering approach uses a single
 1019 file that contains the pre-computed geometry
 1020 describing the model. This file is accessed using
 1021 memory mapping implemented in the Windows™ and
 1022 Linux operating systems. Figure 12 outlines the nec-
 1023 essary function calls for memory mapping the geo-
 1024 metric representation of the vasculature for the
 1025 Windows™ operating system. This has two major
 1026 advantages. First, the file can be randomly accessed.
 1027 File caching is handled entirely by the operating sys-
 1028 tem. This approach utilizes the hardware capabilities
 1029 of the memory management unit (MMU) within the
 1030 CPU. Second, due to the file caching capabilities of
 1031 the operating system, close-up views can be rendered
 1032 at comparatively high frame rates. If the geometry
 1033 that is needed for a close-up view fits into the file
 1034 cache of the computer, no hard disk access is
 1035 necessary making the rendering relatively fast. The
 1036 out-of-core rendering mode has the advantage of

handling data sets larger than the available main 1037
 memory space. 1038

REFERENCES

- 1039
- ¹Andujar, C., C. Saona-Vazquez, I. Navazo and P. Brunet. 1040
 Integrating occlusion culling and levels of details through 1041
 hardly-visible sets. *Computer Graphics Forum* 19(3), 2000. 1042
- ²Bartz, D., M. Meißner, and T. Hüttner OpenGL-assisted 1043
 occlusion culling for large polygonal models. *Computers 1044*
Graphics 23(5):667–679, 1999. 1045
- ³Cohen-Or, D., Y. Chrysanthou, C. T. Silva, and F. Du- 1046
 rand. A survey of visibility for walkthrough applications. 1047
IEEE Trans. Visualization Computer Graphics 9(3):412– 1048
 431, 2003. 1049
- ⁴Deussen, O., C. Colditz, M. Stamminger, and G. Drettakis 1050
 Interactive Visualization of Complex Plant Ecosystems. 1051
IEEE Visualization 2002:219–226, 2002. 1052
- ⁵El-Sana, J and A. Varshney. Generalized view-dependent 1053
 simplification. In P. Brunet & R. Scopigno (eds.) *Computer 1054*

- 1055 Graphics Forum (Eurographics '99), Vol. 18(3). The Eurographics Association and Blackwell Publishers: 83–94, 1999. 1116
- 1056 6El-Sana J., Sokolovsky, C. T. Silva. 2001 Integrated 1117
- 1057 occlusion culling with view-dependent rendering. *IEEE Visualization* 2001:371–378. 1118
- 1058 7El-Sana, J. and E. Bachmat Optimized view-dependent 1119
- 1059 rendering for large polygonal datasets. *IEEE Visualization* 1120
- 1060 2002:77–84, 2002: IEEE Computer Society. 1121
- 1061 8Felkel, P, A. L. Fuhrmann, A. Kanitsar, and R. Wegen- 1122
- 1062 kittel. Surface reconstruction of the branching vessels for 1123
- 1063 augmented reality aided surgery, BIOSIGNAL, Vol. 1, 1124
- 1064 VUTIUM Press: 252–254, 2002. 1125
- 1065 9Gerig, G., T. Koller, G. Széhely, C. Brechbühler, O. Kü- 1126
- 1066 bler. Symbolic description of 3-D structures applied to 1127
- 1067 cerebral vessel tree obtained from MR angiography vol- 1128
- 1068 ume data, *Information Processing in Medical Imaging*, 1129
- 1069 Springer, LNCS: 94–111, 1993. 1130
- 1070 10Greene, N. Hierarchical polygon tiling with coverage 1131
- 1071 masks. *ACM SIGGRAPH* 1996:65–74, 1996. 1132
- 1072 11Gumhold, S. Splatting Illuminated Ellipsoids with Depth 1133
- 1073 Correction. In *Proceedings of 8th International Fall* 1134
- 1074 *Workshop on Vision, Modelling and Visualization* 2003: 1135
- 1075 245–252, 2003. 1136
- 1076 12Hahn, H. K., B. Preim, D. Selle, and H. O. Peitgen Visu- 1137
- 1077 alization and interaction techniques for the exploration of 1138
- 1078 vascular structures. In *IEEE Visualization* 2001:395–402, 1139
- 1079 2001: IEEE, IEEE Computer Society Press. 1140
- 1080 13Hoppe, H. Progressive meshes. *ACM SIGGRAPH* 1141
- 1081 1996:99–108, 1996. 1142
- 1082 14Hoppe, H. Efficient implementation of progressive meshes. 1143
- 1083 *Computers Graphics* 22(1):27–36, 1998: ISSN 0097–8493. 1144
- 1084 15Kaimovitz, B., Y. Lanir, and G. S. Kassab Large-Scale 3-D 1145
- 1085 Geometric Reconstruction of the Porcine Coronary Arte- 1146
- 1086 rial Vasculature Based on Detailed Anatomical Data. *Am. 1147*
- 1087 *Biomed. Eng.* 33(11):1517–1535, 2005. 1148
- 1088 16Kassab, G. S., C. A. Rider, N. J. Tang, and Y. C. Fung 1149
- 1089 Morphometry of pig coronary arterial trees. *Am. J. Physiol.* 1150
- 1090 265(Heart Circ. Physiol. 34):H350–H365, 1993. 1151
- 1091 17Klosowski, J. T. and C. T. Silva Efficient conservative 1152
- 1092 visibility culling using the prioritized-layered projection 1153
- 1093 algorithm. *IEEE Trans. Visualization Computer Graphics* 1154
- 1094 7(4):365–379, 2001. 1155
- 1095 18Linsen, L., B. J. Karis, E. G. McPherson, and B. Hamann 1156
- 1096 Tree Growth Visualization. *J. WSCG* 13:81–88, 2005. 1157
- 1097 19Masutani, Y., K. Masamune and T. Dohi. Region-grow- 1158
- 1098 ing-based feature extraction algorithm for tree-like objects. 1159
- 1099 Visualization in Biomedical Computing, Springer, LNCS: 1160
- 1100 161–171, 1996. 1161
- 1101 20Nordsletten, D. A., S. Blackett, M. D. Bentley, E. L. Rit- 1162
- 1102 man, and N. P. Smith Structural Morphology of Renal 1163
- 1103 Vasculature. *Am. J. Physiol. Heart Circ. Physiol.* 1164
- 1104 291:H296–H309, 2006. 1165
- 1105 21Oeltze, S., B. Preim. Visualization of Anatomic Tree 1166
- 1106 Structures with Convolution Surfaces, IEEE/Eurographics 1167
- 1107 Symposium on Visualization: 311–320, 2004. 1168
- 1108 22Oeltze, S. and B. Preim Visualization of Vasculature with 1169
- 1109 Convolution Surfaces: Method, Validation and Evaluation. 1170
- 1110 *IEEE Trans. Medical Imaging* 24(4):540–548, 2005. 1171
- 1111 23Pajarola, R. FastMesh: efficient view-dependent meshing. 1172
- 1112 In B. Werner (ed.) *Proceedings of the ninth Pacific* 1173
- 1113 *Conference on Computer Graphics and Applications* 1174
- 1114 (PACIFIC GRAPHICS-01). IEEE Computer Society, Los 1175
- 1115 Alamitos, CA, Oct. 16–18: 22–30, 2001. 1176
- 24Puig, A., D. Tost, and I. Navazo An interactive cerebral 1177
- blood vessel exploration system. *IEEE Visualization* 1178
- 97:443–446, 1997. 1179
- 25Reina, G. and T. Ertl. Hardware-Accelerated Glyphs for 1180
- Mono- and Dipoles in Molecular Dynamics Visualization, 1181
- Proceedings of EUROGRAPHICS – IEEE VGTC Sym-* 1182
- posium on Visualization* 2005: 177–182, 2005. 1183
- 26Ritman, E. L. Micro-computed tomography-current status 1184
- and developments. *Annu. Rev. Biomed. Eng.* 6:185– 1185
2082004. 1186
- 27Robb, R. A. The Biomedical Imaging Resource at Mayo 1187
- Clinic. Guest Editorial. *IEEE Trans. Med. Imaging* 1188
- 20(9):854–867, 2001. 1189
- 28Robb, R. A. and C. Barillot Interactive display and anal- 1190
- ysis of 3-D medical images. *IEEE Trans Med Imaging* 1191
- 8(3):217–226, 1989. 1192
- 29Robb, R. A., D. Hanson, R. A. Karwoski, A. G. Larson, 1193
- E. L. Workman, and M. C. Stacy ANALYZE: a compre- 1194
- hensive, operator-interactive software package for multi- 1195
- dimensional medical image display and analysis. 1196
- Computerized Med Imaging Graphics* 13:433–454, 1989. 1197
- 30Segal, M. and K. Akeley. The OpenGL® Graphics System: 1198
- A Specification (Version 2.0–October 22, 2004), [http://](http://www.opengl.org/documentation/specs/version2.0/glspec20.pdf) 1199
- [www.opengl.org/documentation/specs/version2.0/](http://www.opengl.org/documentation/specs/version2.0/glspec20.pdf) 1200
- [glspec20.pdf](http://www.opengl.org/documentation/specs/version2.0/glspec20.pdf). 1201
- 31Shaffer, E. and M. Garland. Efficient Adaptive Simplifi- 1202
- cation of Massive Meshes. In *IEEE Visualization* 2001: 1203
- 127–134, 2001. 1204
- 32Spaan, J. A., R. ter Wee, J. W. van Teeffelen, G. Streekstra, 1205
- M. Siebes, C. Kolyva, H. Vink, D. S. Fokkema, and E. 1206
- VanBavel. Visualization of intramural coronary vascula- 1207
- ture by an imaging cryomicrotome suggests compartmen- 1208
- talization of myocardial perfusion areas. *Med. Biol. Eng. 1209*
- Comput.* 43(4):431–435, 2005. 1210
- 33Stoll, C., S. Gumhold, and H. Seidel Visualization with 1211
- stylized line primitives. *IEEE Visualization* 2005:695–702, 1212
2005. 1213
- 34Woo, M., J. Neider, and T. Davis. OpenGL Programming 1214
- Guide. Addison Wesley, 3rd edn, 2003. 1215
- 35Xia, J., J. El-Sana, and A. Varshney Adaptive real-time 1216
- level-of-detail-based rendering for polygonal models. 1217
- IEEE Trans. Visualization Computer Graphics* 3(2):171– 1218
- 183, 1997. 1219
- 36Yoon, S. E., B. Salomon, and D. Manocha. Interactive 1220
- view-dependent rendering with conservative occlusion 1221
- culling in complex environments. *IEEE Visualization* 2003 1222
- Proceedings*: 163–170, 2003. 1223
- 37Yoon, S. E., B. Salomon, R. Gayle, and D. Manocha. 1224
- Quick-VDR: Interactive View-Dependent Rendering of 1225
- Massive Models. In *IEEE Visualization* 2004 *Proceedings*: 1226
- 131–138, 2004. 1227
- 38Zamir, M. Nonsymmetrical bifurcations in arterial 1228
- branching. *J. General Physiol.* 72(6):837–845, 1978. 1229
- 39Zhang, H., D. Manocha, T. Hudson, and K. E. Hoff III 1230
- Visibility culling using hierarchical occlusion maps. *Com-* 1231
- puter Graphics* 31(Annual Conference Series):77–88, 1997. 1232