

# Parallel Cell Projection Rendering of Adaptive Mesh Refinement Data

Gunther H. Weber<sup>1,2,3</sup>   Martin Öhler<sup>2</sup>   Oliver Kreylos<sup>1,3</sup>   John M. Shalf<sup>3</sup>   E. Wes Bethel<sup>3</sup>  
 Bernd Hamann<sup>1,3</sup>   Gerik Scheuermann<sup>2</sup>

<sup>1</sup> Center for Image Processing and Integrated Computing (CIPIC), Department of Computer Science,  
 One Shields Avenue, University of California, Davis, CA 95616-8562, U.S.A.

<sup>2</sup> AG Graphische Datenverarbeitung und Computergeometrie, FB Informatik, University of Kaiserslautern,  
 Erwin-Schrödinger Straße, D-67653 Kaiserslautern, Germany

<sup>3</sup> Visualization Group, National Energy Research Scientific Computing Center (NERSC),  
 Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, U.S.A.

**Keywords:** volume rendering, adaptive mesh refinement, load balancing, multi-grid methods, parallel rendering, visualization

## ABSTRACT

Adaptive mesh refinement (AMR) is a technique commonly used in numerical simulations of a wide variety of scientific and engineering phenomena or processes. Despite its inherent hierarchical nature, only few visualization algorithms exist that directly use the AMR structure, e.g., for efficient volume rendering schemes. Most existing approaches for AMR data make extensive use of special-purpose graphics hardware. We present an efficient, purely software approach for direct volume rendering of AMR data based on cell-projection utilizing parallel supercomputers or PC clusters. This class of machines is the same as that used to perform the AMR simulations. Thus, the used resources are readily available to users. We discuss a framework that employs this renderer for parallel rendering of AMR data. We introduce and compare several distribution schemes.

## 1 INTRODUCTION

Physical phenomena can vary widely in scale. Large regions in space can exist where a quantity varies only slightly, and thus can be adequately represented at low resolutions. Other regions may require higher resolutions to capture rapid changes. In 1984, AMR was introduced to computational physics by Berger and Olinger [3], aiming to add the ability to adapt mesh resolution during a simulation approach using structured meshes. AMR represents a spatial domain as set of nested structured grids of increasing resolution. Berger and Olinger [3] used a scheme where refining grids can be rotated with respect to a parent level. A modified version of their algorithm was later published by Berger and Colella [2], where all refining grids are axis-aligned with respect to the parent level. AMR has become increasingly popular, also outside the computational physics community. Today, it is used in a large variety of applications. For example, Bryan [4] used the technique to simulate astrophysical phenomena using a hybrid approach combining AMR grids and particles.

Based on an efficient software cell-projection volume renderer, we have developed a framework for parallel volume rendering of AMR data. Even though cell-projection [18] was introduced primarily for rendering unstructured meshes, it also leads to efficient implementations for structured meshes. Our method partitions an AMR hierarchy using a k-d tree [1]. This partition is view-independent and computed offline in a preprocessing step. We

have developed and compared several partition strategies which we briefly summarize.

**Uniform root-level subdivision** ignores the hierarchical nature of AMR data and partitions a root level into blocks of constant size. Refined cells are handled during rendering by recursive descending into finer levels.

**Weighted root-level subdivision** partitions a root level into blocks at approximately constant computational cost. The AMR hierarchy is only considered to compute weights. Locations for subdivision are chosen independently from boundaries of refining grids. During rendering refining grids are handled by descending recursively.

**Homogeneous subdivision** subdivides AMR levels recursively until each part only covers one grid of a given level, i.e., until it corresponds to a region represented at constant resolution. The resulting grid parts are distributed evenly among processors.

**Weighted homogeneous subdivision** partitions AMR levels in the same way as homogeneous subdivision. The computational cost for rendering a constant-resolution region is estimated and associated with that region as its *weight*. Grid parts are distributed among processor such that the sum of associated weights is approximately the same for all processors.

Our framework supports rapid development and testing of new distribution strategies and volume rendering techniques.

## 2 RELATED WORK

Initial work in AMR visualization focused on converting AMR data to suitable conventional representations and visualizing these. Norman *et al.* [23] described a method that visualizes AMR data using standard toolkits. Their method converts an AMR hierarchy into an unstructured grid composed of hexahedral cells. The resulting unstructured grid is then used for visualization utilizing standard algorithms. Observing that by converting AMR data to an unstructured mesh, its main advantage, the implicit definition of grid connectivity, is lost, Norman *et al.* extend VTK to handle AMR as first-class data structure. Max [21] described sorting schemes for cells for volume rendering and described their application to AMR data. Ma [17] described parallel rendering of structured AMR data resulting from simulations using the PARAMESH framework [20]. He described two approaches for volume rendering of AMR data. One method resamples a hierarchy on an uniform grid at the finest resolution. The resulting grid is evenly subdivided and each part




Figure 1: AMR hierarchy consisting of five grids and three levels. Level boundaries are shown as bold lines.

rendered on an individual processor. A second method preserves the AMR structure.

Weber *et al.* [27] presented two volume rendering schemes for Berger-Colella AMR data. One scheme is a hardware-accelerated renderer for previewing; the other scheme allows progressive refinement rendering of AMR data based on cell projection [18]. Weber *et al.* [28] described a method to extract isosurfaces from AMR data. To avoid re-sampling, they interpret locations of cell-centered data values as vertices of a dual grid. Resulting gaps between hierarchy levels are filled via a generic stitching scheme. Weber *et al.* [29] discussed using dual-grids and stitch-cells to define a consistent interpolation scheme for high-quality volume rendering of Berger-Colella AMR data. Kreylos *et al.* [12] described a framework that partitions a Berger-Colella AMR hierarchy partitioned in blocks of constant resolution using a k-d tree. Resulting blocks are distributed among processors and rendered using either a texture-based hardware-accelerated approach or a software-based cell-projection renderer. Ligocki *et al.* [16] described *Chombo-Vis*<sup>1</sup>, a framework for the visualization of hierarchical computations using AMR.

Kähler and Hege [8, 9] introduced a scheme to partition Berger-Colella AMR data in blocks of constant resolution. Aiming to minimize the number of generated constant-resolution blocks they utilize a heuristic based on assumptions concerning the placement of refining grids by an AMR simulation. Kähler *et al.* [10] developed a method that uses AMR hierarchies for rendering sparse volumetric data. Given a transfer function, this method computes a transfer-function-specific AMR hierarchy for a volume data set and renders it using the algorithm of Kähler and Hege [8, 9]. Kähler *et al.* [7] used a set of existing tools to render results of a simulation of a forming star. By specifying a transfer function and a range of isovalues Park *et al.* [24] produced volume-rendered images of AMR data based on hierarchical splatting, see Laur and Hanrahan [13]. Their method converts an AMR hierarchy to a k-d-tree structure consisting of blocks of constant resolution which are rendered back-to-front using hierarchical splatting.

### 3 AMR DATA FORMAT

Figure 1 shows a simple 2D AMR hierarchy produced by the Berger-Colella method. The basic building block of a  $d$ -dimensional Berger-Colella AMR hierarchy is an axis-aligned, structured

<sup>1</sup>Joint effort of the Applied Numerical Algorithms Group and of the Visualization Group at LBNL. See <http://seesar.lbl.gov/anag/chombo/chombovis.html> for further details.

rectilinear grid. Considering the 3D case, each grid  $g$  consists of hexahedral cells and is positioned by specifying its local origin  $\mathbf{o}_g$ . AMR typically uses a *cell-centered* data format, *i.e.*, dependent function values are associated with cells/cell centers. Data values are stored in arrays as location and connectivity can be inferred from the regular grid structure.

An AMR hierarchy consists of several levels  $\Lambda_l$  comprising one or multiple grids. All grids in the same level have the same cell size. A hierarchy's *root level*  $\Lambda_0$  is the coarsest level. Each level  $\Lambda_l$  may be refined by a finer level  $\Lambda_{l+1}$ . A grid of a refined level is referred to as a *coarse grid* and a grid of a refining level as a *fine grid*. A *refinement ratio*  $r$  specifies how many fine grid cells fit into a coarse grid cell, considering all axis-directions. This value is always a positive integer. A refining grid refines an entire level  $\Lambda_l$ , *i.e.*, it is completely contained in the region covered by that level but not necessarily in the region covered by a single grid of that level. Each refining grid can only refine complete grid cells of the parent level, *i.e.*, it must start and end at the boundaries of grid cells of the parent level. The Berger-Colella scheme [2] requires the existence of a layer with a width of at least one grid cell between a refining grid and the boundary of the refined level.

## 4 DESIGN CONSIDERATIONS

One can differentiate volume rendering methods by their underlying illumination models (*i.e.*, the “optical properties” of transfer functions) and by their operation in *image* or *object space*. Two illumination models are widely used in volume rendering: The absorption and emission light model, described, for example by Max [22] and the Phong-based light model by Levoy [14]. We chose the absorption and emission light model as it leads to efficient implementations. Within cells, we use constant interpolation, *i.e.*, the sample value located at the cell center is assigned to all positions within the cell. This allows an exact evaluation of the light-model and preserves the AMR hierarchy in rendered images. To achieve more efficiency, we chose orthographic projection over perspective projection.

Image-space-based algorithms, including the commonly used ray casting algorithm, see Sabella [26], operate on pixels in screen space as “computational units,” *i.e.*, they perform computations on a per-pixel basis. Object-space-based methods, like cell projection, see Ma and Crockett [18], operate on 3D grid cells. Parallelizing volume rendering can be done in image space or in object space, see Crockett [5]. Image-space parallelization subdivides the image plane to distribute computing among multiple processors. Each processor renders a subset of pixels in an image. Object-space parallelization subdivides the domain of a data set and assigns grid cells to processors. We chose object-space based parallelization, as the hierarchical nature of AMR data facilitates efficient subdivision of the grids. We chose cell-projection as an object-space-based rendering methods, as it leads to an elegant implementation of subdivision of the domain. Furthermore, using cell-projection eases reaction to changes in resolution, *i.e.*, it is possible to render finer grids at a higher resolution.

For implementation of the parallel renderer we chose the Message Passing Interface (MPI) library over the Parallel Virtual Machine (PVM) framework. MPI is commonly used in AMR simulations, thus making our framework more compatible with other applications, including numerical simulation. Furthermore, MPI is a de facto standard for parallel supercomputers. Vendor-specific adaptations for different architectures exist, supporting the utilization of specific hardware optimizations by linking to a vendor-provided library. Instead of adopting the classic master-slave model, we chose a symmetric implementation to avoid communication bottlenecks. Each processor computes the complete distribution of grid parts and selects a subset based on its index. However,

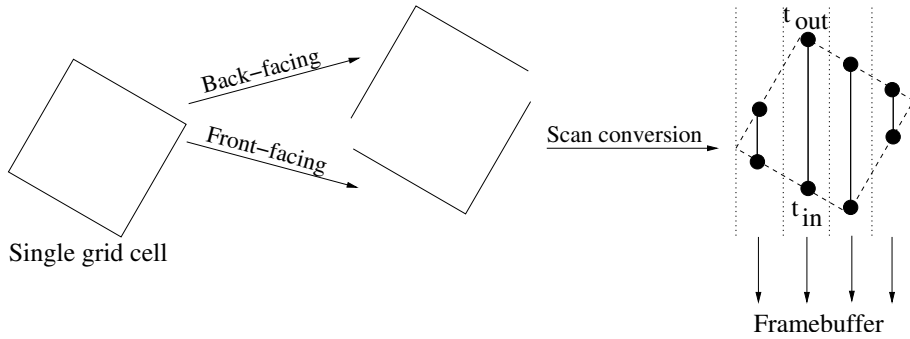


Figure 2: Cell-projection process.

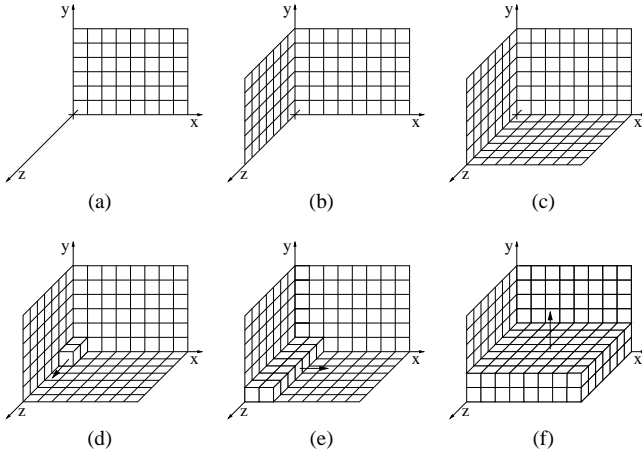


Figure 3: Rendering order of grid cells — all components of  $\mathbf{tv}$  being positive. First, all back-facing faces of the first layer of cells in each direction are rendered (a) – (c). Second, all cells are rendered. The order in which axes are handled (first- $z$ —then- $x$ —then- $y$  order) is arbitrary. Only the order according to which cells are handled along an axis is important.

we are using a binary-tree image compositing scheme that pairs processors in each compositing step. In each step, one processor of each pair receives an intermediate partial image from its “neighbor” and performs a compositing operation. The final composited image resides in the buffer of processor zero.

## 5 EFFICIENT SOFTWARE-BASED CELL-PROJECTION

### 5.1 Overview

Cell projection [18] is an object-space-based volume rendering method, similar in nature to ray casting. Both methods trace rays through a volume, accumulating light along the path of a ray. Ray casting operates on a per-pixel basis, using one ray for each pixel. Cell-projection-based methods construct “ray segments” for cells and merge them with existing ray segments.

Usually, a priority queue is maintained for each pixel collecting all ray segments contributing to that pixel. Figure 2 shows the fundamental idea of cell projection. Boundary faces of all cells are divided into three groups, *front-facing* faces (with normals

directed toward the viewer) and *back-facing* faces (with normals directed away from the viewer). First, the back-facing faces are scan-converted into a buffer. For each pixel influenced by the cell, this buffer holds a depth corresponding to an exit parameter value, called  $t_{out}$ , along the ray. Second, the front-facing faces are scan-converted. For each generated pixel, the depth corresponding to the entry parameter value, called  $t_{in}$ , along the ray is computed. The entry parameter value  $t_{in}$  and corresponding scalar value are read from the buffer, and the ray segment reaching from  $t_{in}$  to  $t_{out}$  is constructed. Usually, this ray segment is then inserted into the ray-segment queue of the corresponding pixel and merged with adjacent ray segments in that queue.

When cells are sorted using, for example, the scheme of Max [21] and rendered in back-to-front or front-to-back order, the queue for collecting ray segments is not necessary. Newly generated ray segments are always adjacent to already computed ray segments and can be composited directly in the frame buffer. Another advantage of this method is that it allows us to avoid duplicate scan conversion of a cell’s boundary faces. When rendering unsorted cells, back-facing and front-facing faces must be rendered to determine correct ray-segment length. In contrast, when rendering presorted cells, it is sufficient to render the front-facing faces of a cell. All back-facing faces are already rendered as front-facing faces of cells “behind” the current cell.

### 5.2 Cell Sorting and Front-face Determination

For AMR grids it is simple to determine correct back-to front cell order. For orthographic projections, rendering order can be determined based on view direction. Cells are enumerated by three nested loops, one loop for each axis. The order according to which axes are handled is arbitrary. Along each axis, cells must be handled in correct order. For each loop this order can be determined based on the sign of the component of the vector  $\mathbf{tv}$  (a vector directed toward the viewer), according to the axis handled by the loop. If it is positive, cells are enumerated in ascending axis direction. If it is negative, cells are enumerated in descending axis direction. If it is zero, an arbitrary choice is made.

Before rendering cells and generating ray segments, all cell faces lying on the up to three back-facing boundary-faces of the overall AMR grid that are not view-perpendicular must be scan-converted. These are the back-facing faces of cells that do not lie in front of any grid cell. Figures 3 (a)–(c) illustrate the procedure for a choice of  $\mathbf{tv}$  where all components are positive. If one component of  $\mathbf{tv}$  is zero, the corresponding face is perpendicular to the viewing direction and discarded. Subsequently, the front facing faces of all cells are scan-converted. Ray segments are generated and composited in the frame buffer. Figure 3 shows the order of scan-conversion




Figure 4: Determining front- and back-facing faces. (a) The vector  $\mathbf{tv}$  is perpendicular to the viewing plane, pointing toward the viewer. (b) Face numbering used in Table 1.

Face #	Front-facing	Back-facing	Perpendicular
0	$tv_z < 0$	$tv_z > 0$	$tv_z = 0$
1	$tv_x > 0$	$tv_x < 0$	$tv_x = 0$
2	$tv_z > 0$	$tv_z < 0$	$tv_z = 0$
3	$tv_x < 0$	$tv_x > 0$	$tv_x = 0$
4	$tv_y < 0$	$tv_y > 0$	$tv_y = 0$
5	$tv_y > 0$	$tv_y < 0$	$tv_y = 0$

Table 1: Criteria used for checking whether a cell face is front-facing, back-facing, or should not be rendered.

used for boundary faces and cells when all components of  $\mathbf{tv}$  are positive.

When dealing with more general grid cells, each boundary face must be checked individually whether it is back-facing or front-facing. This step can be done, for example, by using the scalar product between face normal and  $\mathbf{tv}$ . For axis-aligned rectilinear grids, this test can be performed based on the viewing direction. Figure 4 shows the numbering of the faces of a cell of a rectilinear grid. Table 1 lists the criteria used to determine whether a face is front- or back-facing. Front- and back-facing faces are the same for each cell in all grids. It is sufficient to determine front-facing faces once.

### 5.3 Boundary Face Scan-conversion

We render cell boundary faces using a modified version of the polygon scan-conversion algorithm developed by Gordon *et al.* [6] that is based on a method developed by Kaufman [11]. Before rendering a polygon, its vertices are projected onto the viewing plane, and point coordinates are rounded to integers. During the scan-conversion process it is assumed that coordinates are specified counter-clockwise. Gordon *et al.*'s method starts by determining "critical points" of a polygon, *i.e.*, vertices that constitute a local minimum or are first of a set of vertices that together form a local minimum in  $y$ -direction. Boundary faces of rectilinear cells are convex quadrilaterals and have only one minimum. If two adjacent vertices share the same value for the  $y$ -coordinate, *i.e.*, if they are connected by a horizontal line segment, we consider the vertex with the lower index to be the critical point. During the determination of critical points we also detect polygons that span one pixel in  $y$ -direction and discard them.

The algorithm starts by inserting the left minimum and right maximum edge originating from the critical point vertex into an active-edge table (AET). During the scan-conversion process this data structure holds the left and right edge intersecting the current scan-line. For general polygons, considered by Gordon *et al.* [6] this is an array, as the scan-line can intersect several polygons.




Figure 5: Scan-converted polygon illustrating rules used to determine whether a pixel belongs to the polygon.

Our scan-converter is optimized for convex quadrilaterals and only stores two pointers to AET elements since a convex quadrilateral intersects a scan-line only twice, except for horizontal boundary lines coinciding with a scan-line. For each scan-line,  $x$ -coordinates and depth on the left and the right side of the polygon are calculated by linear interpolation. Depth information is not rounded, as exact values are needed for the determination of ray-segment lengths. If a scan-line consists of only one pixel of the polygon it is discarded; otherwise, depth values are computed for all pixels between the  $x$ -coordinates by linear interpolation. Ray segments are created by reading the previous depth value and applying the illumination model. If a scan line coincides with the end of an edge, the corresponding pointer referring to the AET is replaced with its successor until that turns down.

When generating images with a cell-projection method it is important that rasterized polygons sharing an edge do not overlap. Thus, special care must be taken at polygon boundaries. We use these rules:

- R1. Integer intersection points of a polygon edge with a scan-line belong to a polygon, if they lie on its left edge. If they lie on its right edge, they do not belong to the polygon.
- R2. Non-integer intersection points of a polygon edge with a scan-line are rounded down. The corresponding pixel belongs to a polygon if it lies on its right edge. If it lies on its left edge, it does not belong to the polygon.
- R3. If a pixel corresponds to intersection points on the left and right edges of a polygon it lies outside the polygon.

The white center polygon in Figure 5 illustrates these rules: The pixel at the lower-left corner of the polygon has an integer intersection point and lies on its left edge. Consequently, according to R1, it belongs to the polygon. According to R1, the pixel at the upper right corner of the polygon does not belong to the polygon. Considering R2, all pixels with non-integer intersection points on the left and lower polygon edge do not belong to the polygon. (They belong to the neighboring polygon.) All pixels bordered by the upper

and right polygon edges do belong to the polygon. The pixel at the upper-left corner lies on the left and the right edges of the polygon and does not belong to the polygon (R3). During scan-conversion, we maintain a list of all positions modified, *i.e.*, covered by a cell. This list is used in the compositing scheme, see Section 7, and to speed up clearing the frame buffer by only erasing pixels modified during rendering.

## 5.4 Ray-segment Generation

We use constant interpolation in individual cells, *i.e.*, the sample value associated with a cell is assigned to all positions in the cell. Consequently, all points in a cell have the same optical properties, *i.e.*, emission color and opacity. It is possible to solve the differential equations for light absorption and emission analytically in a cell and obtain “correct” opacity and emission values for a ray segment intersecting the cell. Each ray segment in a cell is characterized by an entry parameter value  $t_{in}$ , *i.e.*, the distance from the viewing plane at which a ray “enters” the cell measured along the ray, and an exit parameter value  $t_{out}$ , *i.e.*, the distance from the viewing plane at which the ray “exits” the cell. The value of  $t_{in}$  is obtained by scan-converting the front-facing faces of a cell. The value of  $t_{out}$  is read from the frame buffer containing the results from scan-converting the front faces of cells (behind the current cell) that coincide with the back-facing faces of the current cell. Emission color and opacity are defined by the cell’s associated scalar value via a transfer function.

Max [22] described the the absorption and emission light model using an extinction coefficient  $\tau_{Cell}$  instead of opacity  $\alpha_{Cell}$  per unit length. This extinction coefficient can be obtained from the opacity as

$$\tau_{Cell} = -\ln(1 - \alpha_{Cell}) . \quad (1)$$

Using Max’s equations from [22] we obtain the transparency from  $t_{in}$  to an arbitrary parameter value along the ray in the cell as

$$T_{Cell}(s) = \exp\left(-\int_{t_{in}}^s \tau_{Cell} dk\right) = (1 - \alpha_{Cell})^{s-t_{in}} \quad (2)$$

The opacity of a ray segment is

$$\alpha_{Seg} = 1 - T_{Cell}(t_{out}) = 1 - (1 - \alpha_{Cell})^{t_{out}-t_{in}} . \quad (3)$$

Furthermore, the color (intensity) of a ray segment can be computed as

$$C_{Seg} = \int_{t_{in}}^{t_{out}} C_{Cell} \tau_{Cell} T(s) ds = C_{Cell} \alpha_{Seg} , \quad (4)$$

using again equations from [22]. Combining contributions of individual ray segments is equivalent to compositing pixels with a color  $C_{Seg}$  and an opacity  $\alpha_{Seg}$  using pre-multiplied alpha values, see Porter and Duff [25], *i.e.*,

$$\begin{aligned} \alpha_{Combined} &= 1 - (1 - \alpha_{Seg,front})(1 - \alpha_{Seg,back}) \quad \text{and} \\ C_{Combined} &= \alpha_{Combined} ((1 - \alpha_{Seg,front}) C_{Seg,back} + C_{Seg,front}) . \end{aligned} \quad (5)$$

## 6 PARTITIONING AND LOAD-BALANCING

### 6.1 Overview

Our method stores a domain partition as a k-d tree [1]. A k-d tree is a generalization of a binary search-tree to arbitrary dimensions. Each level partitions a domain in two regions along an axis-perpendicular plane. The “left” sub tree corresponds to points in space whose coordinates in partition direction have values smaller than or equal to the partition position. The “right” sub tree corresponds to points whose coordinates in partition direction have




Figure 6: Uniform subdivision of root level into equal-sized blocks.

values larger than the partition position. The subdivision direction is usually alternated between the three coordinate axes in the 3D case. We skip subdivision directions, when no “sensible” subdivision position along that direction exists. When using an object-space-based subdivision for parallelizing volume rendering, regions must be rendered in correct order. Using k-d trees makes it possible to determine this order simply. At each node of the k-d tree, the domain is subdivided along an axis-perpendicular plane. The compositing order can be determined by considering the component of  $t_{\mathbf{v}}$  corresponding to the partition direction. If it is positive the left sub tree must be rendered first; if it is negative the right sub tree must be rendered first; and if it is zero both sub trees can be rendered in arbitrary order. We assume that subdivision schemes are view-independent. It is sufficient to compute a k-d tree subdivision once per data set. We compute subdivisions offline in a pre-processing step. To assign regions of the domain, *i.e.*, leaves of the k-d tree, to individual processors they are numbered in rendering order. We assign a set of sequentially adjacent regions to each processor.

### 6.2 Uniform Root-level Subdivision

Given an AMR hierarchy, this scheme constructs a k-d tree with a user-specified number of levels. Each node of the tree splits its associated region into two parts of nearly equal size, *i.e.*, number of cells. Figure 6 shows uniform subdivision of the AMR hierarchy from Figure 1.

Since uniform subdivision ignores grid boundaries, refined cells must be handled during rendering. We implemented a solution method based on recursively descending the hierarchy. While rendering a data set, a test is performed for each grid cell checking whether it is refined by the next finer level. If a finer level exists, the  $r^3$  ( $r$  being the refinement ratio) refining grid cells are rendered instead of the current coarser cell. The correct rendering order of refining cell is determined using the criteria described in Section 5. Each refining cell is checked recursively to determine potential further refinement.

### 6.3 Weighted Root-level Subdivision

Similarly to uniform subdivision, this scheme ignores grid boundaries of an AMR hierarchy during subdivision. The goal of this approach is to obtain a subdivision of a given AMR hierarchy into regions that will imply approximately equal computational cost. With each region we associate an estimate of computational cost for rendering, used as a weight. A subdivision plane is chosen using a greedy method as follows: Initially, the subdivision plane is placed




Figure 7: Weighted subdivision of root level, ignoring grid boundaries.

CPU type	$c_0$	$c_1$	$c_2$
1.0 GHz AMD Athlon	1.00	0.60	0.50
1.2 Ghz AMD Athlon	1.00	0.65	0.54
1.4 GHz AMD Athlon	1.00	0.71	0.58
2.4 Ghz Intel Xeon	1.00	0.63	0.54
375 MHz IBM Power 3	1	0.77	0.71

Table 2: Constants for weighted distribution for different processors.

in the middle of the current domain. Weights are computed for the two subdomains. If both subdomains have equal weight, the subdivision plane has optimal position and the algorithm terminates. Otherwise, the plane is moved into the subdomain with the larger associated weight. Moving the plane in this way decreases computational cost for that subdomain while increasing computational cost for the other one. We calculate the weight difference before and after moving the plane, and continue moving the plane as long as it decreases the weight difference. Our algorithm terminates when moving the plane increases the difference instead of decreasing it, or the partition plane would reach the border of a subdomain.

We estimate computational cost of a subdomain based on the number of cells in it. We must also consider the fact that rendering refined cells is computationally more expensive than rendering unrefined cells. Therefore, we assign a weight of one to unrefined cells of the root level, *i.e.*,  $c_0 = 1$ . Based on an application specific benchmark it is possible to determine relative weights for refined cells. Table 2 shows weights for an AMR hierarchy consisting of three levels. The constants specify the times necessary to render a single cell of a given AMR level with respect to rendering times for a single cell of the root level. These constants are measured by rendering a cell of the appropriate level from a viewing direction of  $t_v = (1, 1, 1)$ . Viewing a cell in this direction no cell faces are axis-perpendicular. A maximum number of faces must be rendered and the “footprint” of the cell on the screen has maximum size. The associated weight  $w$  of a subdomain is

$$w = \sum_{l=0}^{\text{\#Level}} n_l c_l, \quad (6)$$

where  $n_l$  is the number of level  $l$  cells. This sum is computed by recursively descending in the hierarchy. Figure 7 shows weighted subdivision of the root level for the AMR hierarchy from Figure 1, using relative cell weights of  $c_0 = 1$ ,  $c_1 = 0.75$  and  $c_2 = 0.7$ .




Figure 8: Homogeneous subdivision of AMR hierarchy.

## 6.4 Homogeneous Subdivision

Both sub-division strategies discussed so far ignore the hierarchical nature of AMR data during subdivision. Only the impact on computational cost for rendering a grid part is considered when using weighted root-level subdivision. Resulting regions usually encompass several grids of the original hierarchy, leading to data duplication and poor memory utilization. By considering grid boundaries during the subdivision step, it is possible to partition an AMR hierarchy into “homogeneous” blocks, *i.e.*, blocks represented at constant resolution. Each block contains only cells from one grid of the original AMR hierarchy. This property allows us to avoid data duplication. Due to the homogeneous nature of blocks, it is possible to render them efficiently, avoiding tests for refinement of individual cells and recursion.

Subdivision of an AMR data set uses only information about the hierarchical structure of AMR data. Actual data values for individual grids do not need to be loaded, and subdivision can be performed on a single machine, requiring only a small amount of memory. We construct the k-d tree level by level. For level  $l$ , we locate all leaf nodes of the current k-d tree that correspond to grids of level  $l - 1$ . Each of these leaves is replaced by a k-d tree that is constructed as follows: We determine all grids of level  $l$  that overlap the leaf region, *i.e.*, the region associated with the current leaf. Since grids may only partially overlap the leaf region, they are clipped against the leaf region to obtain the grid part contained in the leaf region. Along the current subdivision direction, we store every position in subdivision direction where a refining grid starts or ends.

After sorting the resulting list and removing duplicate elements, we choose the middle element of the resulting list as subdivision position. (If the list contains an even number of elements, we choose the smaller of the two middle elements as subdivision position. If the list is empty, we skip the corresponding subdivision direction.) For each of the two regions associated with a subtree of the current leaf, we find all grids that overlap that region and clip them against the boundaries of that region. Alternating between the three axis directions, we repeat this process recursively until all leaf regions are homogeneous and overlap only a single grid. For the root level, we start with an empty tree that covers the complete domain and construct a k-d tree analogously to creating the tree for a leaf.

By terminating k-d tree construction after a user-specified fixed level-number it is possible to render only a part of an AMR hierarchy. Figure 8 shows the results of homogeneous subdivision of the AMR hierarchy from Figure 1. Grid parts are numbered in back-to-front order and distributed among processors. Each processor loads the complete partition information and renders nearly the same number of sequentially numbered grid parts.

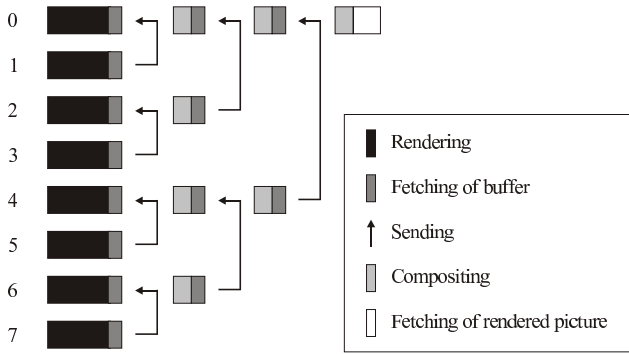


Figure 9: Parallel compositing scheme.

## 6.5 Weighted Homogeneous Subdivision

Weighted homogeneous subdivision uses the same k-d tree subdivision as homogeneous subdivision. Instead of distributing resulting grid parts evenly among processors, an estimate of computational rendering cost is used as a weight for each leaf of the k-d tree. This weight is obtained by multiplying the number of grid cells in the leaf region by the weight of a single cell of the appropriate level. The weight of a single cell is the same as the one used in weighted root-level subdivision, see Section 6.3. Regions are distributed among processors using a greedy method. To achieve nearly equal processor utilization, each processor needs to render regions with a total weight of  $w_{\text{ideal}} = \frac{w_{\text{Total}}}{71 \text{ Processors}}$ , where  $w_{\text{Total}}$  is the computational cost for rendering the complete AMR hierarchy, *i.e.*, the sum of all weights of all leaves of the k-d tree. Each processor has an assigned set of sequentially adjacent leaves. Processor  $p$  selects its assigned regions as follows: If  $k$  is the last leaf rendered by processor  $p - 1$ , processor  $p$  adds the remainder of that region, *i.e.*, the part that was not rendered by the previous processor, to its “assignment list.” (An exception to this rule applies to the first processor. It does not need to render any partial regions.) Starting with region  $k + 1$ , processor  $p$  adds regions to its assignment list until the weight of the current region exceeds the difference  $w_{\text{ideal}} - w_{\text{curr}}$ . During each step,  $w_{\text{curr}}$  denotes the sum of weights of all regions already assigned to processor  $p$ .

To achieve a more uniform distribution of weights, this region is subdivided as follows: First, we choose the direction perpendicular to the plane consisting of the least number of cells as partition direction. We divide the difference  $w_{\text{ideal}} - w_{\text{curr}}$  by the weight of a slice in partition direction (*i.e.*, the number of cells in the slice multiplied by the cell weight of the appropriate level). The result is rounded to obtain the number of slices rendered by the current processor. The remaining slices are rendered by the next processor. (An exception to this rule applies to the last processor which renders all remaining regions.)

Each processor computes assignments for all processors. This avoids the need for waiting for the previous processor to finish its own assignment computation. The index  $k$  of the last region rendered by the previous processor and a potential remainder of a subdivided region are determined locally. Performing this computation in parallel, avoids added time for communication between processors.

## 7 COMPOSITING

When all regions are rendered, the partial images are composited. Compositing is done by using alpha blending/compositing of the partial images, see Porter and Duff [25]. The compositing scheme



Figure 11: Processor utilization for uniform root-level subdivision.

is illustrated in Figure 9. In the first step, each odd processor sends its partial image to its lower-indexed neighbor processor that performs the compositing operation. In each subsequent step  $i$  only those processors that composited an image in the previous step are considered, *i.e.*, processors with an index of the form  $k2^i$ . Each processor having an associated odd value of  $k$  sends its intermediate partial image to processor  $(k - 1)2^i$  which performs the next compositing operation. Because regions are assigned to processors in back-to-front order, each processor can composite the partial image received from the other processor “over” the region in its own buffer. At the end of the compositing process, the final image resides in the buffer of processor zero. When transferring partial images between processors for compositing, we only transmit pixels that have been altered during rendering. We do this by transferring position, color and alpha value for each altered pixel. In the “fetch buffer” stage this representation is converted to a bitmap.

## 8 RESULTS

Figure 10 shows the last time step of the “Argon Bubble” data set. We used this data, which is courtesy of Center for Computational Sciences and Engineering (CCSE), Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, California, for testing our distribution strategies. It is the result from the last time-step in a simulation of a shock wave passing through an Argon bubble surrounded by another gas. The visualized scalar field is gas density. This simulation is stored in AMR format using a hierarchy consisting of 885 grids in three levels. All grids in total consist of 1401504 grid cells. Homogeneous subdivision of the AMR hierarchy yields 6002 grid regions. Figure 10(a) shows the grid structure. Figure 10(b) shows the final volume-rendered image.

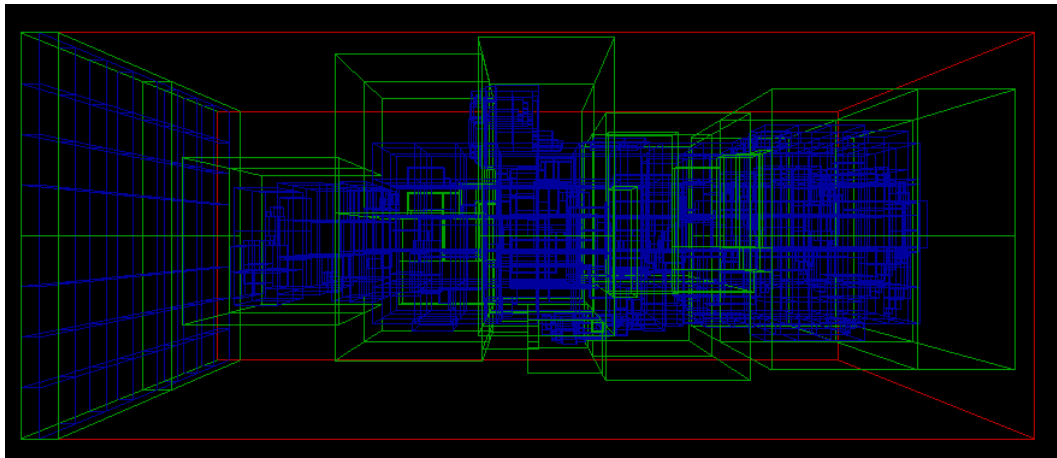
We tested distribution strategies on the following machines:

**Linux Cluster** This configuration is a Linux cluster consisting of four 1.2 GHz Dual-Athlon machines connected via a Gigabit network. For measurements with four processors, we used a single CPU on each machine. For measurements with eight processors, we used both CPUs on each machine. Each machine has 512 MB main memory.


**Shared-memory machine** This is a PC-based server equipped with two 2.4 GHz Intel Xeon CPUs using hyper-threading to obtain four “virtual” CPUs. The used machine has a total memory of 2 GByte RAM. We used a version of MPICH that supports the shared memory environment on that machine.

**IBM SP2** Seaborg is a 10 Teraflop IBM SP RS/6000 located at NERSC’s high-performance computing facility. It consists of 416 NightHawk II nodes. Each node contains 16 IBM Power3+ processors running at 375 MHz and 16–64 GBytes of shared memory. The nodes are interconnected using dual 150 Megabyte/s SP/“GX BusColony” switch adaptors forming a fat-tree topology. We used IBM’s native MPI implementation.

Figures 11 – 14 show processor utilization for rendering on a four-processor Linux cluster. As expected, uniform subdivision



(a)



(b)

Figure 10: (a) Grid structure of “Argon Bubble.” The hierarchy consists of 885 grids in three levels with a root grid of  $80 \times 32 \times 32$  cells. (b) Volume-rendered image of “Argon Bubble.” (Data set courtesy of Center for Computational Sciences and Engineering (CCSE), Ernest Orlando Lawrence Berkeley National Laboratory, Berkeley, California)




Figure 12: Processor utilization for weighted root-level subdivision.




Figure 13: Processor utilization for homogeneous subdivision.

achieves an uneven utilization of processors. Weighted root-level subdivision achieves a comparatively even processor utilization, but it requires longer rendering times than subdivisions working on homogeneous grid parts. This behavior is due to the overhead by recursively descending into the hierarchy. Recursive descend also causes non-local memory access pattern resulting in poor cache utilization. Working only on data of a single grid and avoiding overhead due to data inhomogeneity, homogeneous subdivision performs better than weighted root-level subdivision, even though computational cost is not as evenly distributed. Weighted homoge-

neous subdivision resolves this problem and achieves good rendering speed while near-uniformly utilizing all processors.

Tables 3 – 5 show rendering times and speedups for rendering on a Linux cluster and a shared memory machine. Speedups are measured with respect to rendering the data on a single processor using homogeneous subdivision. As expected, weighted homogeneous subdivision leads to best results of all considered subdivision schemes. We note that times vary between subsequent runs of our framework and timings are only accurate within approximately one second. Considering these facts, the speedup achieved by weighted



No. of CPUs	Weighted Root-level		Homogeneous		Weighted Homogeneous	
	Time [s]	Speedup	Time [s]	Speedup	Time [s]	Speedup
1	142.10	1.00	125.45	1.00	124.98	1.00
4	39.87	3.56	36.25	3.46	33.83	3.69
8	21.67	6.55	19.81	6.33	17.89	6.98
16	11.61	12.23	12.93	9.70	8.79	14.21
32	7.88	18.03	7.31	17.16	6.24	20.02
64	5.42	26.21	6.22	20.16	2.97	42.08
128	3.09	45.98	3.83	32.75	1.98	63.12
256	2.07	68.64	1.48	84.76	1.53	81.68
512	1.66	85.60	1.27	98.77	1.37	91.22

Table 6: Rendering times on IBM SP2.




Figure 14: Processor utilization for weighted homogeneous subdivision.

Subdivision Strategy	Time [s]	Speedup
Uniform	14.76	2.59
Weighted Root-level	14.60	2.62
Homogeneous	12.16	3.14
Weighted Homogeneous	11.23	3.40

Table 5: Rendering times on shared-memory machine.

Subdivision Strategy	Time [s]	Speedup
Uniform	14.90	2.56
Weighted Root-level	14.67	2.61
Homogeneous	12.35	3.10
Weighted Homogeneous	11.95	3.20

Table 3: Rendering times on Linux Cluster using four CPUs.

Subdivision Strategy	Time [s]	Speedup
Uniform	7.83	4.89
Weighted Root-level	7.50	5.10
Homogeneous	7.10	5.39
Weighted Homogeneous	6.21	6.16

Table 4: Rendering times on Linux Cluster using eight CPUs.




Figure 15: Speedup as function of number of processors (IBM SP2).

homogeneous subdivision is satisfactory.

Table 6 shows rendering times on an IBM SP2. These measurements are of “strong scaling” behavior whereby the problem size remains fixed as the number of processors is increased. This typically results in less flattering scaling efficiency than if the problem size was scaled to be proportional to the number of processors as is the case for “weak scaling” studies. Timing granularity for large-scale parallel applications is typically on the order of approximately one second. Thus, results utilizing more than 128 processors on that system have a lower degree of confidence than the smaller tests. Starting with 256 processors, homogeneous subdivision surprisingly performs better than weighted homogeneous subdivision. The difference in the performance of the models for the very large scale runs is less than the timing granularity, so we only have limited confidence that these effects are actually real rather than being timing artifacts. However, that being said, we believe these timings are consistent with the observation that the time required for assigning regions to processors is higher for weighted homogeneous subdivision. While rendering time on each processor decreases for a larger number of processors, the time spent computing the subdivision becomes the dominant computational cost. It is also possible that the granularity of work that can be assigned becomes large

compared to the total amount of work that is assigned to each processor — offering less benefit to these fine-grained optimizations.

## 9 CONCLUSIONS AND FUTURE WORK

We have implemented and compared several distribution strategies for direct volume rendering of AMR hierarchies. Homogeneous subdivision supports efficient rendering of AMR data for different classes of machines. It allows us to avoid data duplication and employ a wide variety of rendering schemes. Homogenizing an AMR hierarchy has also been used for a variety of hardware-accelerated methods for volume-rendering AMR data. While weighted homogeneous subdivision of the domain results in near-uniform processor utilization, we plan to improve the approximation of relative cell weights. It may be beneficial to use view-dependent weights. We intend to consider inhomogeneous PC clusters consisting of machines with varying processor speed and develop subdivision/distribution methods that take these differences in machine performance into account. Furthermore, we plan to develop a communication-less subdivision strategy that avoids the need for

each processor to compute assignments for all other processors, resulting in less overhead and better scalability. During compositing, a major portion of time is spent on sending and receiving partial images. We plan to reduce this overhead by encoding partial images more efficiently using, for example, run-length encoding. We intend to implement binary-swap compositing [19] and compare its results to ours.

## ACKNOWLEDGMENTS

We thank Hartmut Sprengart from the Centrum für Produktionstechnik (CCK)/Lehrstuhl für Fertigungstechnik und Betriebsorganisation (FBK) for permission to use their shared-memory Intel-Xeon machine. This work was supported by the Stiftung Rheinland-Pfalz für Innovation; by the Director, Office of Science, of the U. S. Department of Energy under Contract No. DE-AC03-76SF00098; the National Science Foundation under contract ACI 9624034 (CAREER Award), through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251, and through the National Partnership for Advanced Computational Infrastructure (NPACI); the National Institute of Mental Health and the National Science Foundation under contract NIMH 2 P20 MH60975-06A2; and the Lawrence Berkeley National Laboratory.

We thank the members of the AG Graphische Datenverarbeitung und Computergeometrie at the Department of Computer Science at the University of Kaiserslautern, Germany, the Visualization and Graphics Research Group at the Center for Image Processing and Integrated Computing (CIPIC) at the University of California, Davis, and the Visualization Group at the Lawrence Berkeley National Laboratory.

## REFERENCES

- [1] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, September 1975.
- [2] Marsha Berger and Phillip Colella. Local adaptive mesh refinement for shock hydrodynamics. *Journal of Computational Physics*, 82:64–84, May 1989. Lawrence Livermore National Laboratory, Technical Report No. UCLRL-97196.
- [3] Marsha Berger and Joseph Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *Journal of Computational Physics*, 53:484–512, March 1984.
- [4] Greg L. Bryan. Fluids in the universe: Adaptive mesh refinement in cosmology. *Computing in Science and Engineering*, 1(2):46–53, March/April 1999.
- [5] Thomas W. Crockett. An introduction to parallel rendering. *Parallel Computing*, 23(7):819–843, July 1997.
- [6] Dan Gordon, Michael A Peterson, and R. Anthony Reynolds. Fast polygon scan conversion with medical applications. *IEEE Computer Graphics and Applications*, 14(6):20–27, November 1994.
- [7] Ralf Kähler, Donna Cox, Robert Patterson, Stuart Levy, Hans-Christian Hege, and Tom Abel. Rendering the first star in the universe – a case study. In: Robert J. Moorhead, Markus Gross, and Kenneth I. Joy, editors, *IEEE Visualization 2002*, pages 537–540, IEEE, IEEE Computer Society Press, Los Alamitos, California, 2002.
- [8] Ralf Kähler and Hans-Christian Hege. Interactive volume rendering of adaptive mesh refinement data. Technical Report ZR-01-30, Zuse Institut Berlin (ZIB), Berlin, Germany, 2001. Appeared in *The Visual Computer* [9]. Available as <ftp://ftp.zib.de/pub/zib-publications/reports/ZR-01-30.pdf>.
- [9] Ralf Kähler and Hans-Christian Hege. Texture-based volume rendering of adaptive mesh refinement data. *The Visual Computer*, 18(8):481–492, 2002.
- [10] Ralf Kähler, Mark Simon, and Hans-Christian Hege. Fast volume rendering of sparse high-resolution datasets using adaptive mesh refinement hierarchies. Technical Report ZR-01-25, Zuse Institut Berlin (ZIB), Berlin, Germany, 2001. To appear in *IEEE Transactions on Visualization and Computer Graphics*. Available as <ftp://ftp.zib.de/pub/zib-publications/reports/ZR-01-25.pdf>.
- [11] Arie Kaufman. Efficient algorithms for scan-converting 3d polygons. *Computers & Graphics*, 12(2):213–219, 1988.
- [12] Oliver Kreylos, Gunther H. Weber, E. Wes Bethel, John M. Shalf, Bernd Hamann, and Kenneth I. Joy. Remote interactive direct volume rendering of amr data. Technical Report LBNL 49954, Lawrence Berkeley National Laboratory, 2002.
- [13] David Laur and Pat Hanrahan. Hierarchical splatting: A progressive refinement algorithm for volume rendering. *Computer Graphics (Proceedings of ACM SIGGRAPH 91)*, 25(4):285–288, July 1991.
- [14] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, May 1988. (See also corrigendum [15, 30]).
- [15] Marc Levoy. Letter to the editor: Error in volume rendering paper was in exposition only. *IEEE Computer Graphics and Applications*, 20(4):6–6, July/August 2000.
- [16] Terry J. Ligocki, Brian Van Straalen, John M. Shalf, Gunther H. Weber, and Bernd Hamann. A framework for visualizing hierarchical computations. In: Gerald Farin, Bernd Hamann, and Hans Hagen, editors, *Hierarchical and Geometrical Methods in Scientific Visualization*, pages 197–204. Springer Verlag, Heidelberg, Germany, January 2003.
- [17] Kwan-Liu Ma. Parallel rendering of 3D AMR data on the SGI/Cray T3E. In: *Proceedings of Frontiers '99 the Seventh Symposium on the Frontiers of Massively Parallel Computation*, pages 138–145, IEEE, IEEE Computer Society Press, Los Alamitos, California, February 1999.
- [18] Kwan-Liu Ma and Thomas W. Crockett. A scalable parallel cell-projection volume rendering algorithm for three-dimensional unstructured data. In: James Painter, Gordon Stoll, and Kwan-Liu Ma, editors, *IEEE Parallel Rendering Symposium*, pages 95–104, IEEE, IEEE Computer Society Press, Los Alamitos, California, November 1997.
- [19] Kwan-Liu Ma, James S. Painter, Charles D. Hansen, and Michael F. Krogh. Parallel volume rendering using binary-swap composition. *IEEE Computer Graphics and Applications*, 14(2):59–67, July 1994.
- [20] Peter MacNeice, Kevin M. Olson, Clark Mobarry, Rosalinda de Fainchtein, and Charles Packer. Paramesh: A parallel adaptive mesh refinement community toolkit. *Computer Physics Communications*, 126(3):330–354, April 2000.
- [21] Nelson L. Max. Sorting for polyhedron compositing. In: Hans Hagen, Heinrich Müller, and Gregory M. Nielson, editors, *Focus on Scientific Visualization*, pages 259–268. Springer-Verlag, New York, New York, 1993.
- [22] Nelson L. Max. Optical models for volume rendering. *IEEE Transactions on Computer Graphics*, 1(2):99–108, 1995.
- [23] Michael L. Norman, John M. Shalf, Stuart Levy, and Greg Daus. Diving deep: Data management and visualization strategies for adaptive mesh refinement simulations. *Computing in Science and Engineering*, 1(4):36–47, July/August 1999.
- [24] Sanghun Park, Chandrajit Bajaj, and Vinay Siddavanahalli. Case study: Interactive rendering of adaptive mesh refinement data. In: Robert J. Moorhead, Markus Gross, and Kenneth I. Joy, editors, *IEEE Visualization 2002*, pages 521–524, IEEE, IEEE Computer Society Press, Los Alamitos, California, 2002.
- [25] Thomas Porter and Tom Duff. Compositing digital images. *Computer Graphics (Proceedings of ACM SIGGRAPH 84)*, 18(3):253–259, July 1984.
- [26] Paolo Sabella. A rendering algorithm for visualizing 3D scalar fields. *Computer Graphics (Proceedings of ACM SIGGRAPH 88)*, 22(4):51–58, 1988.
- [27] Gunther H. Weber, Hans Hagen, Bernd Hamann, Kenneth I. Joy, Terry J. Ligocki, Kwan-Liu Ma, and John M. Shalf. Visualization of adaptive mesh refinement data. In: Robert F. Erbacher, Philip C. Chen, Jonathan C. Roberts, Craig M. Wittenbrink, and Matti Groehn, editors, *Proceedings of the SPIE (Visual Data Exploration and Analysis VIII, San Jose, CA, USA, Jan 22–23)*, volume 4302, pages 121–132, SPIE, SPIE – The International Society for Optical Engineering, Bellingham, WA, January 2001.
- [28] Gunther H. Weber, Oliver Kreylos, Terry J. Ligocki, John M. Shalf, Hans Hagen, Bernd Hamann, and Kenneth I. Joy. Extraction of crack-free isosurfaces from adaptive mesh refinement data. In: David S. Ebert, Jean M. Favre, and Ronny Peikert, editors, *Proceedings of the Joint EUROGRAPHICS and IEEE TCVG Symposium on Visualization, Ascona, Switzerland, May 28–31, 2001*, pages 25–34, 335, Springer Verlag, Wien, Austria, May 2001.
- [29] Gunther H. Weber, Oliver Kreylos, Terry J. Ligocki, John M. Shalf, Hans Hagen, Bernd Hamann, Kenneth I. Joy, and Kwan-Liu Ma. High-quality volume rendering of adaptive mesh refinement data. In: Thomas Ertl, Bernd Girod, Günther Greiner, Heinrich Niemann, and Hans-Peter Seidel, editors, *Vision, Modeling, and Visualization 2001*, pages 121–128, 522. Akademische Verlagsgesellschaft Aka GmbH, Berlin, Germany and IOS Press BV, Amsterdam, Netherlands, November 2001.
- [30] Craig Wittenbrink, Tom Malzbender, and Mike Goss. Letter to the editor: Interpolation for volume rendering. *IEEE Computer Graphics and Applications*, 20(5):6–6, September/October 2000.